Australian Government

Department of Defence
Defence Science and
Technology Organisation

The State-of-the-Art in Ship Detection in Synthetic Aperture Radar Imagery.

D. J. Crisp

DSTO-RR-0272

20040915 059

Australian Government

Department of Defence

Defence Science and
Technology Organisation

The State-of-the-Art in Ship Detection in Synthetic Aperture Radar Imagery.

D. J. Crisp

Intelligence, Surveillance and Reconnaissance Division
Information Sciences Laboratory

DSTO–RR–0272

ABSTRACT

This report is a review of the publicly available literature on algorithms for detecting ships in Synthetic Aperture Radar (SAR) imagery, with the aim of recommending algorithms for inclusion in the Analysts' Detection Support System (ADSS). The ADSS is a software system being developed at DSTO for the automatic detection of targets in SAR imagery. We outline the motivations, theory and justifications of the various approaches so that educated comparisons and evaluations can be made. Most current research on ship detection is based on detecting the ship itself rather than its wake. Detection of ship wakes is not addressed in any depth. Ship detection systems generally consist of several stages: land masking; preprocessing; prescreening; and discrimination. We consider each of these stages in turn and discuss the various algorithms which have been used. A basic system for low to medium resolution single channel SAR imagery is recommended and suggestions for improvements and alternatives are made.

APPROVED FOR PUBLIC RELEASE

AQ F04-11-1296

APPROVED FOR PUBLIC RELEASE

The State-of-the-Art in Ship Detection in Synthetic Aperture Radar Imagery.

EXECUTIVE SUMMARY

This report is a review of the publicly available literature on algorithms for detecting ships in Synthetic Aperture Radar (SAR) imagery, with the aim of recommending algorithms for inclusion in the Analysts' Detection Support System (ADSS). Our intention is not just to summarise the current state of affairs in the literature. We also aim to outline the motivations, theory and justifications of the various approaches so that educated comparisons and evaluations can be made.

The driving requirement at DSTO for ship detection in SAR imagery is wide area surveillance of the oceans surrounding Australia and especially to the North where tropical atmospheric conditions hinder surveillance with electro-optic sensors. Modern SAR sensors can generate large amounts of data in a short period of time and there is an obvious need for automatic detection. This is especially true in the context of ship surveillance where much of the imagery contains only open ocean. The ADSS is a software system being developed at DSTO for the automatic detection of targets in SAR imagery. It was designed with the intention of reducing imagery analysts' workloads in wide area surveillance operations by directing their attention to areas in the imagery of likely interest.

The launch of the Seasat SAR system in 1978, verified that ships could be imaged with space-based SAR systems. It also revealed an unexpected richness of oceanographic phenomena. Seasat was followed by several other space-based SAR systems including ERS1, ERS2, Radarsat and SIR-C/XSAR. These systems have made high quality SAR imagery of the ocean readily available and operational surveillance of ocean ship traffic is now viable. Several near real-time operational ship monitoring systems using SAR imagery are now in place or are currently being developed. The prime examples are the Ocean Monitoring Workstation (OMW), the Alaska SAR Demonstration (AKDEMO) system, the European Community Joint Research Centre (JRC) system, and the Qinetiq's MaST system. The Norwegian Defence Research Establishment (NDRE) is an important player too and was behind the system Eldhuset reported on in 1996.

Given the long history and ongoing interest there is an extensive literature on algorithms for ship detection in SAR imagery. While many papers begin with a brief review of the literature relevant to their application, only one general review has been found. That review, by Fingas and Brown, notes the lack of previous reviews but is itself limited. This report aims to fill that gap.

There are two fundamentally different means of detecting ships in SAR imagery: detection of the ship target itself and detection of the ship wake. While it is clearly desirable to make use of both possibilities (and many ship detection systems do), in this report, we only focus on algorithms which detect the ship itself. This decision was made in order to limit the size of the report. It is anticipated that the detection of ship wakes will be the topic of a companion report. As a justification of this stance, we note that most current research on ship detection is based on detecting the ship rather than its wake. Among the reasons for this are: ship backscatter is robust and less dependent on sea state; stationary

ships have no wake; wakes are often not visible and especially at large radar incidence angles; and wake detection is more computationally expensive.

We also comment that there is a bias in the open literature towards space based rather than airborne SAR systems. In theory this should make little difference to detection algorithms since the imaging principles for both platforms are the same. However, we note that space-based SAR imagery is generally single channel and lower resolution. Thus most ship detection algorithms in the literature are aimed at locating bright point-like targets in single channel imagery. Space based SARs also tend to have lower incidence angles which makes the detection problem harder. We add however, that the next generation of space-based SARs offer multi-channel polarimetric data and consequently there is a shift towards research on algorithms which make use of multi-channel data.

Past research efforts on automatic target detection in SAR imagery have clearly demonstrated that no single detection algorithm will produce satisfactory results and a hierarchical system of algorithms is needed. Thus ship detection systems generally consist of several stages: land masking; preprocessing; prescreening; and discrimination. In this report, we consider each of these stages in turn and discuss the various algorithms which have been used. Recommendations on algorithms for the ADSS are made in the final section.

There are two main difficulties in making recommendations. First, different types of imagery require different algorithms. Resolution is of prime importance here but other factors such as polarisation and image processing are relevant too. Second, there is a lack of rigorous performance evaluation in the literature. Often the imagery is visually ground-truthed and some algorithms are very poorly tested with results only being presented for a handfull of images. This makes comparison of the various approaches difficult. Given this, our first recommendation is to:

- specify a surveillance scenario;

- decide on appropriate radar parameters and geometry;

- set up well ground-truthed benchmark tests;

- compare the different detection algorithms.

The only complete systems to have extensive and convincing publicly available reports of their performance are the OMW, AKDEMO, SUMO and Eldhuset systems. On this basis we recommend the following basic system for low to medium resolution single channel SAR imagery:

- land masking using a shore-line database with a buffer zone included;

- a simple moving window adaptive threshold algorithm (with the two parameter CFAR being the preferred option);

- clustering of detected pixels;

- discrimination of false alarms using elementary target measurements;

- human supervision and discrimination of the final detections;

We note that this system can be applied to higher resolution imagery if an initial down sampling step is included. Once this basic system is in place, improvements and alternatives can be considered. We discuss the details of this basic system and make suggestions for improvements and alternatives.

We believe that while many of the suggested refinements and alternatives may result in improved performance, the most significant gains are likely to come from future research on polarimetric imagery and better false alarm discrimination. Much of this research will be based on Envisat ASAR and Radarsat-2 data.

Author

David J. Crisp
Intelligence, Surveillance and Reconnaissance Division

David Crisp graduated from the University of Adelaide in 1987 with a B.Sc. (Hons) in Mathematics and completed his Ph.D. at the same institution in 1993. For the three years following that he held a postdoctoral research position in Mathematics at the Flinders University of South Australia. In 1997 he commenced employment as a Postdoctoral Research Fellow with the Cooperative Research Centre for Sensor, Signal and Information Processing (CSSIP). At CSSIP he worked in the Pattern Recognition Group on the application of machine learning techniques to real world problems. In September 1999 he joined the Surveillance Systems Division of the Australian Defence Science and Technology Organisation (DSTO). Since joining DSTO he has been a Research Scientist in the Image Analysis and Exploitation Group and his research has been focused on the automated detection of targets in wide area surveillance synthetic aperture radar imagery.

Contents

Figures

Tables

Glossary

ADSS: Analysts' Detection Support System

AKDEMO: Alaska SAR Demonstration

DSTO: Defence Science & Technology Organisation

CFAR: Constant False Alarm Rate

ESA: European Space Agency

FAR: False Alarm Rate

InSAR: Interferometric SAR

JRC: European Community Joint Research Centre

MaST: Maritime Surveillance Tool

OMW: Ocean Monitoring Workstation

PD: Probability of Detection

pdf: Probability Density Function

PFA: Probability of False Alarm

RCS: Radar Cross Section

SAR: Synthetic Aperture Radar

SUMO: Search for Unidentified Maritime Objects

SLC: Single Look Complex

1 Introduction

This report is a review of the published literature on algorithms for detecting ships in Synthetic Aperture Radar (SAR) imagery with the aim of describing algorithms for potential inclusion in the Analysts' Detection Support System (ADSS). Our intention is not just to summarise the current state of affairs in the literature. We also aim to outline the motivations, theory and justifications of the various results and approaches so that educated comparisons and evaluations can be made.

Modern SAR sensors can generate large amounts of data in a short period of time and there is an obvious need for automatic detection of targets of interest. This is especially true in the context of ship surveillance where much of the imagery contains only open ocean. Automation also has the potential to guarantee consistency and predictability. At DSTO, a software system called the ADSS is being developed for the automatic detection of targets in SAR imagery, see [101] and [100]. The ADSS was designed with the intention of reducing imagery analysts' workload in wide area surveillance operations by directing their attention to areas in the imagery of likely interest. The design and development of the ADSS is ongoing. The driving requirement for ship detection with SAR imagery at DSTO is wide area surveillance of the oceans surrounding Australia and especially the oceans to the North where tropical atmospheric conditions hinder surveillance with electro-optic sensors.

While many papers on ship detection begin with a brief review of the relevant literature, only one general review has been found. That review, by Fingas and Brown [37], notes the lack of previous reviews. While the Fingas and Brown review is worthwhile, it is clearly limited. Its scope includes optical sensors and it cites less than 30 papers in total. It is also very brief being only 7 pages long. Of the many other papers on ship detection, good overviews may be found in Yeremy *et al.* [142], Wackerman *et al.* [135], Henschel *et al.* [47], Eldhuset [32] and Wahl [138]. A snap shot of the state-of-the-art in the year 2000, can be found in the Fall issue of Backscatter [11]. This issue contains amongst other items, summaries of workshops on "Ship Detection for Fisheries Monitoring" [59], "Ship Detection with Satellite-based Sensors" [132] and "Ship Detection from Airborne Platforms" [36]. Another snap shot of activity was provided by the symposium on "Emerging and Coastal Applications of wide swath SAR" held in March 1999 at the Johns Hopkins University Applied Physics Laboratory. The proceedings were published in Johns Hopkins APL Technical Digest, Volume 21, Number 1, 2000. Among the papers of interest are [44], [56], [94], [89], [82] and [126]. Many other meetings on ship detection have been held but the proceedings are not readily available.

We do not attempt to completely cover the topic of ship detection in SAR imagery and two restrictions have been placed on the scope of this report. First, there are two fundamentally different means of detecting ships in SAR imagery: detection of the ship target itself and detection of the ship wake. While it is clearly desirable to make use of both possibilities (and many ship detection systems do), we will only focus on algorithms which detect the ship itself. This decision was made in order to limit the size of the report. It is anticipated that the detection of ship wakes will be the topic of a companion report. As a justification of this stance, we note that most current research on ship detection is based on detecting the ship rather than its wake. Among the reasons for this are: ship

backscatter is robust and less dependent on sea state; stationary ships have no wake; wakes are often not visible and especially at large radar incidence angles; and wake detection is more computationally expensive.

The second limit to the scope of this report relates to the expertise in the group at DSTO behind the development of the ADSS. The skills of this group are founded on rigorous approaches using mathematical modelling and statistical hypothesis testing. Therefore, less rigorous techniques such as artificial neural networks, genetic programming, etc., will not be covered in detail. We will indicate such possibilities with some representative papers but we will not delve into the details. This perspective agrees with that of Oliver and Quegan who note in the introduction to their book [86] that *ad hoc* methods rarely prove to be robust and rigorous approaches are preferred.

We also note that not all work on ship detection is being published in the open literature. Given the nature of the topic, it is to be expected that many reports are of defence significance and therefore not publicly available. It is also the case that some of the research is commercially sensitive and therefore only published as internal reports or not at all. The prime example of this is Qinetiq's ship detection system, MaST. The only reference we found to MaST was the web page [96]. Similarly, not all the documentation for Satlantic's ship detection system, the OMW, is readily available. For example, Yeremy *et al.* [142], cite the internal Satlantic report [95], by Plache and Henschel. Where possible such internal reports have been obtained (although [95] could not be). Also, some work has not been published at all. For instance, the bibliography of the IGARSS'01 conference lists the papers: "Comparison and Validation of Different Algorithms for (Merchant and Fishing) Ship Detection with Spaceborne SAR Data" by P-L. Frison *et al.*; and "Improvement and Validation of Ship Detection by SAR Interferometry" by A. Arnaud; yet these papers did not appear.

Past research efforts on automatic target detection in SAR imagery have clearly demonstrated that no single detection algorithm will produce satisfactory results and a hierarchical system of algorithms is needed. Thus ship detection systems generally consist of several stages: land masking; preprocessing; prescreening; and discrimination. In this report, we consider each of these stages in turn and discuss the various algorithms which have been used. This discussion forms the content of Sections 3 to 8. In order to put the discussion in context, we first review, in Section 2.1, the topic of SAR imagery and its nature and properties in ocean surveillance applications. Following the discussion of the various detection stages, we turn in Section 9 to a description of the ship detection systems as a whole. In this section we also discuss the various research groups and organisations behind the systems and indicate their future research directions where possible. In the final section of the report we make recommendations on ship detection algorithms for the ADSS.

2 SAR imagery for ship detection

2.1 SAR imagery

The aim of a synthetic aperture radar (SAR) is to form an image which represents the radar reflectivity of the scene being imaged. The radar reflectivity is expressed in terms of a quantity known as the *radar cross section* (RCS) and denoted by σ. For various reasons, including the comparison of SAR systems with different resolutions, it is usual to remove the area dependence from RCS measurements. Thus the output of a SAR system is usually the *normalised radar cross section* NRCS which is the RCS per unit area illuminated. The NRCS is also known as the *radar backscattering coefficient* and is almost universally denoted by σ^0.

While other radar systems can also be used to form images, the distinguishing feature of SAR is that it can achieve a high resolution in the azimuth direction as well as the range direction. This is done mathematically through the construction of a synthetic aperture. The details are not relevant here. There are several different SAR image formation algorithms corresponding to the different modes of operation — stripmap, spotlight, scan, etc. However, we do not propose to go into the technical details of SAR image formation here. While it is possible that an understanding of the SAR image formation process can help with detector design, mostly that is not the case and all that is required is a general understanding of the nature and properties of SAR imagery. Good background references on SAR image formation are, for instance, [18] and [86].

An important factor to take into account when discussing SAR imagery is the platform. A SAR system can be spaceborne or airborne. Spaceborne SARs are usually simple due to the complexity and expense of deployment. Thus multi-frequency, multi-polarisation SARs are more common on airborne platforms and the image quality tends to be better. On the other hand, spaceborne SARs provide wider spatial coverage and regular revisits. Generally, spaceborne SAR is better for wide area surveillance and airborne SAR is better for flexible tasking. Consequently, most ship detection systems have been designed for satellite SAR imagery. However, as Fingas and Brown point out, [37] and [36], much of the work is applicable to either platform because the principles are the same.

Clearly, detection algorithms need to take account of the nature and properties of the imagery they are applied to. In this respect, the most important SAR parameters are

- resolution,

- incidence angle,

- polarisation,

- radar frequency.

Other relevant factors are: pixel spacing, number of looks, swath width, etc. For satellite SAR, the orbital parameters are also of interest: nominal altitude, inclination, sun synchronicity, length of repeat cycle. We will discuss some of these factors and their affect on ship detection below and in the following subsections. However, our intention is only to

cover the most important aspects and to indicate the type of considerations that need to be made.

As just mentioned, a fundamental parameter of SAR imagery is the resolution. Maritime surveillance generally involves a trade off between resolution and coverage. See for instance the article on Radarsat SAR mode selection for marine applications by Vachon and Olsen, [131]. Higher resolution allows for higher probabilities of detection, especially for smaller ships, but it comes at the cost of narrower swath widths and longer revisit times. Thus the highest resolution SAR imagery is not necessarily the best. For example, Kourti *et al.* [61] and also Gower and Skey [44] and [43], recommend the use of Radarsat ScanSAR Narrow imagery for surveillance of fishing vessels. The resolution of such imagery is in the order of 50–100 metres whereas the highest resolution Radarsat Fine mode has a resolution of 8 metres. The gain is that the swath width of ScanSAR Narrow imagery is 305 km compared with 45 km for Fine mode imagery and hence whole fisheries can be covered with a single ScanSAR Narrow image. We also add here Lombardo and Sciotti's remarks that while higher resolution imagery allows for more reliable discrimination and possibly classification of ships, it comes at the cost of greater inhomogeneity and very non-Gaussian statistics which need to be taken into account in detector design. The segmenting detector they use for high resolution imagery is more complex than the one they use for low resolution imagery. However, we emphasise that the choice of resolution depends on the size and nature of the ships to be detected and that small wooden ships may not show up in low resolution imagery.

Due to the geometry and physics of the formation process, SAR imagery has some peculiar properties. They include:

- speckle noise,

- effects from lack of scene coherence,

- azimuth shifting and smearing of moving objects,

- radar shadow,

- foreshortening and layover,

- multiple bounce effects.

While all of these can affect the design and operation of ship detection and classification algorithms, the first three are especially important and we discuss them further.

Speckle is a fundamental property of SAR imagery. Oliver and Quegan, [86], emphasise that speckle is only noise-like rather than an actual noise process. Its cause lies in the scattering of coherent radar waves from rough surfaces. Oliver and Quegan present a detailed explanation of speckle. As is well known, the effect of speckle is to give the observed SAR intensity I a negative exponential distribution

$$p(I) = \frac{1}{\sigma} \exp\left(-\frac{I}{\sigma}\right), \qquad I \geq 0, \tag{1}$$

where σ is the "underlying RCS". As we discuss in the section on prescreening algorithms, the cell averaging CFAR (CA-CFAR) detector was specifically designed for imagery with a

negative exponential distribution. The CA-CFAR is used in many ship detection systems. Often SAR imagery is multi-look processed to reduce the effects of speckle. Multi-look processing involves incoherently averaging several independent measurements of RCS (referred to as looks) and it effectively trades resolution off against RCS accuracy. Again the details are in Oliver and Quegan's book. If L looks are averaged then observed intensity I is known to obey a gamma distribution with order parameter L:

$$p(I) = \frac{1}{\Gamma(L)} \left(\frac{L}{\sigma}\right)^L I^{L-1} \exp\left(-\frac{LI}{\sigma}\right), \qquad I \geq 0, \tag{2}$$

Multi-looking is just one example of how the processing stream affects the statistics of SAR imagery and we repeat that care needs to be taken to ensure algorithms are appropriate for the imagery being used.

Another effect related to the nature of the SAR imaging process is image blurring caused by a lack of scene coherence. One of the assumptions underlying the SAR image formation process, is that the scene being imaged behaves coherently for the duration of the imaging process. Loosely, a scene is coherent if the nature of the scattering mechanisms in the scene do not change. This assumption is clearly violated for dynamic surfaces like wind blown foliage or the ocean surface. A fundamental quantity for gauging such effects is the **coherent integration time** (also sometimes referred to as the exposure time) which is the length of time a scatterer spends in the radar beam. It is given by

$$T \approx \frac{\lambda}{D}\frac{R}{V} \tag{3}$$

where λ is the radar frequency and D is the length of antenna in the azimuth direction. Vachon *et al.*, [130], state that typically T is of the order of 1–10 seconds, although we note that $T = 0.5$s for Radarsat. The **scene coherence time** τ is the time period over which the scene can be considered to retain the coherence necessary for proper SAR image formation. If the scene coherence time is less than the SAR coherent integration time then the azimuth resolution of the SAR image will be degraded. Vachon *et al.* cite from the literature that this loss of resolution is

$$\rho = \rho_1 \sqrt{L^2 + \frac{T^2}{\tau^2}} \tag{4}$$

where ρ_1 is the azimuth resolution for single look imagery without effects from the lack of scene coherence and L is the number of looks. Vachon *et al.* make no comment on the possible values of τ for the ocean surface. However estimates do exist, and Carande, [17], for instance, in experiments with L-band along-track interferometric data from AirSAR, estimates the coherence time of the ocean surface to be in the range from 60 ms to 120 ms which is much less than usual coherent integration times for SAR systems.

The theoretical basis of the SAR imaging principle also assumes that all objects in the scene being imaged are stationary. When this assumption does not hold and targets are in motion some significant effects result. The most important of these is **azimuth image shift**. Azimuth image shift is the effect whereby a target moving with a non-zero range velocity component will not appear in its correct position in the image but will be displaced in the azimuth direction. The formulae for quantifying this displacement are

well known, see [106]. For an airborne SAR, if V is the aircraft ground speed, R the range to the target and u_R the range velocity component of the target motion then the azimuth shift is given by

$$\delta x = -u_R \frac{R}{V}. \tag{5}$$

The factor R/V in this formula is known as the **range velocity ratio**. It is clearly an important factor in the design of SAR sensors and is often quoted in the specifications. In general, the smaller R/V is the better. For satellite SAR, the ground speed V is different to the satellite speed V_p and the formulae becomes

$$\delta x = -u_R \frac{R V_p}{V^2}. \tag{6}$$

Note here that $V_p/V = (r + h)/r$ where r is the radius of the earth and h is the height of the satellite above the earth. Azimuth image shift leads to the curious situation where, for example, a ship may be displaced from its wake in SAR imagery. In this situation, the range component of ship velocity can be inferred from the size of the displacement by inverting either formulae (5) or (6). Note that if the ship orientation is also known then the true velocity can be estimated in this way. The factor R/V in (5) is known as the **range velocity ratio**. It is clearly an important factor in the design of SAR sensors and is often quoted in the specifications. In general, the smaller R/V is the better. Radarsat has a nominal range velocity ratio of 115s. Other effects of target motion in the range direction are range walk and amplitude reduction, see [106] for details.

Target motion in the azimuth direction also affects SAR imagery although not to the same extent as azimuth image shift. It causes blurring in the azimuth direction and hence is called **azimuth defocusing** or **azimuth smearing**. Azimuth smearing is actually the result of second order motion towards the radar due to the combination radar motion and target azimuth motion. For exactly the same reason, a steady acceleration of the target in the range direction also causes azimuth smearing. It follows that ship pitching, rolling and yawing motions and also heave, sway and surge motions will lead to azimuth defocusing. However, nothing has been said about this in the ship detection literature, except for a passing suggestion by Iehara *et al.*, [53] that standard ship detection algorithms may struggle in high seas due to the blurring caused by ship motion. We also mention here the paper by Ouchi *et al.*, [92], which discusses non-uniform azimuth shifts of moving ships and attributes the cause to ship pitching motions. Of course, ocean waves will also be subject to these effects.

SAR image quality issues also need to be considered. Again, it is beyond the scope of this report to delve too deeply here and our intention is only to give a few examples. A partial list includes:

- system noise (measured by the noise equivalent sigma naught σ_{NE}^0),
- image blurring (defocusing),
- geometric distortion,
- system point spread function (PSF) characteristics,
- sidelobes - both temporal and spatial (aspects of the PSF),

- radiometric accuracy,

- polarimetric accuracy,

- azimuth banding (caused by antenna yaw),

- scalloping (horizontal banding due errors in calculating the azimuth beam pattern),

- radio interference.

We cite some examples of these for Radarsat SAR data next.

Radarsat SAR has been extensively used for ship surveillance purposes and its image quality issues are well understood. Some examples from the ship detection literature are as follows. Wackerman *et al.*, [135], note that system noise can be a problem. It causes bright "dots" which their detection algorithm mistakes for ships. They also comment on the presence of bright stripes oriented in the range direction which cause some false detections. They attribute these stripes to the SAR processor used but do not explain exactly what the problem is. Werle, [140], reports the occurrence of azimuth ambiguity patterns. In effect, azimuth ambiguity patterns are "ghosted" replicas of real image features but shifted in azimuth. They are due to higher than desired side lobes which in turn are caused by the low pulse repetition frequency necessary to achieve large incidence angles. Werle gives examples where this leads to ghosting of bright land features over the ocean. Such patterns represent a significant problem to ship detection algorithms since they can be mistaken for real ships. Similarly, temporal sidelobes can result in range displaced false targets, see for instance the work of Schwartz *et al.* in [114]. Vachon *et al.* [126] report image quality problems due to automatic gain control (AGC) errors causing saturation of the analog-to digital converter (ADC). Henschel *et al.* [47] report the need to mask out beam seam/nadir ambiguities. Gower and Skey, [43], report the presence of very visible narrow vertical lines as a result of nadir ambiguities. Rey *et al.*, [102], note that false alarms can be caused by poor SAR processing and heavy saturation in the image. Kourti *et al.*, [61], reports Radarsat ScanSAR errors resulting in bright spots at the edges of the beams. Hawkins *et al.*, [46], report that due to the complexity of combining several elevation beams and the inherent processing issues, it appears that image quality cannot be sustained with the desired radiometric fidelity with ScanSAR. They also mention beam seam/nadir ambiguities and AGC problems. More recently, Aresu and Schwartz, [4], have compared the image quality and radiometric calibration of several different Radarsat ScanSAR processors with automatic ship detection and classification in mind. Amongst other things, they comment on underflow and overflow problems and also on output scaling and range (8 bits or 16 bits) problems. They note in particular that while logarithmic output scaling is useful for oil spill detection, it is not suitable for ship detection. Robertson *et al.*, [105], also considered the effect of using different SAR processors and found a small improvement with an experimental 16 bit processor.

Having noted the work of Werle [140], where the problem of azimuth ambiguity patterns caused by low pulse repetition frequencies (PRFs) is discussed, we add here a comment on the article by Luscombe by Lightstone [76], where it is suggested that using even lower PRFs on Radarsat may be advantageous for ship detection purposes. The idea here is that by tolerating the azimuth ambiguities (in fact they are predictable and hence do not prevent accurate detection) other gains can be made. The proposed mode of operation

allows for wide coverage at relatively high incidence angles both of which are desirable for ship detection purposes.

2.2   SAR imaging systems

Both airborne and spaceborne SAR systems have been used for ship detection. The main space-based systems referred to in the literature are Seasat, ERS-1, SIR-C/X-SAR, ERS-2, Radarsat-1, Envisat and Radarsat-2. Note that Envisat has only recently been launched and Radarsat-2 is not due until late 2005. A brief review of these systems (except Radarsat-2) is given by Raney in [97]. We repeat part of Raney's table of system parameters in Table 1. We have updated the entries for Envisat and added a new entry for Radarsat-2. Sources for these entries are given below. The term **QuadPol** in the table means that all four polarisation channels are measured and hence the radar is fully polarimetric. We discuss polarisation further in Section 6. We stress that the table contains nominal values and is a guideline only. SAR systems often have several operating modes and system parameters such as incidence angle, polarisation, resolution and swath width will vary depending on the mode of operation. We expand on this below. It should also be borne in mind that incidence angle and range resolution are only specified for the mid-swath location and will vary across the swath. Other relevant satellite systems which we mention here in passing are JERS-1 (L-band, HH, 39^o), ALMAZ-1 (S-band, HH, 30–60^o) and the upcoming ALOS PALSAR (L-band, QuadPol, 8–60^o).

The main airborne systems are NASA's AirSAR and CCRS's C/X SAR. The latter is flown on a Convair-580 aircraft and consequently is often referred to as the Convair-580 SAR. We repeat Raney's table of parameters for these systems too in Table 1. However, care is needed here — airborne SAR systems are typically used for technology development as well as applications and so their hardware is frequently changed. The entries for NASA's AirSAR and CCRS's C/X SAR were current in 1995 but are likely to have changed now. Raney comments that AirSAR is the most widely deployed airborne SAR system and its QuadPol capability at three frequencies makes it a rich data source. Raney also observes that the R/V ratio of SIR-C/X-SAR and the airborne SAR systems are much more favourable for imaging the ocean than the other systems. Both AirSAR and the Convair-580 SAR are also capable of along-track (or time difference) interferometry, we discuss this mode of operation and how it can be used for ship detection in Section 7. Note that the parameters of the CCRS's C/X SAR are somewhat similar to those of Radarsat-2. Interestingly, Yeremy *et al.* comment in [143] that the C/X SAR has in fact been used to provide an indication of Radarsat-2's capabilities.

Next we provide more details on the Seasat, ERS-1/2, Radarsat-1/2 and Envisat systems. We begin with Seasat. Seasat was the first spaceborne SAR system. Moreover Seasat was designed specifically for viewing the ocean. Raney, [97], comments that Seasat set the standard for virtually all satellite SAR systems that followed. Precedents for certain image parameters, and in particular resolution and number of looks were established by Seasat and there is little flexibility unless one is willing to sacrifice other parameters such as swath width. While Seasat has been historically important in the development of ship detection algorithms, it only flew for 3 months and is really no longer relevant.

Like Seasat, the ERS-1/2 systems have a fixed viewing geometry and are single channel.

8

SAR sensor	Seasat	ERS-1/2	Radarsat-1	SIR-C/X-SAR
Country	U.S.A.	Europe	Canada	USA/Germany/Italy
Launch date	1978	1991/1995	1995	1994
Lifetime	3 months	2–3 years	5 years	10 days
Wavelength band	L	C	C	L; C; X
Polarisation	HH	VV	HH	L&C - QuadPol; VV
Incidence angle (^{o})	23	23	<20–50	15–55
Range resolution (m)	25	26	10–100	10–30
Azimuth resolution (m)	25	28	9–100	30
Looks	4	6	1–8	~ 4
Swath width (km)	100	100	10–500	15–60
Noise equiv. σ^0 (dB)	-24	-24	-23	-40; -40; <-20
Nominal R/V (s)	115	115	115	33
Exposure time (s)	2.5	0.7	0.5	0.7; 0.16; 0.085

SAR sensor	Envisat	Radarsat-2	AirSAR	C/X SAR
Country	Europe	Canada	USA	Canada
Launch date	2002	2005		
Lifetime		7 years		
Wavelength band	C	C	C; L; P	C; X
Polarisation	Co & CrossPol	QuadPol	QuadPol	C - QuadPol
Incidence angle (^{o})	15–45	10–60	20–60	0–85
Range resolution (m)	30–1000	3–100	7.5 (slant)	6–20
Azimuth resolution (m)	30–1000	3–100	2	<1–10
Looks			4	1–7
Swath width (km)	60–400	20–500	7–13	18–63
Noise equiv. σ^0 (dB)				-40; -30
Nominal R/V (s)			25	45
Exposure time (s)			1.1; 3.7; 9.5	2; 1.2

Table 1: Satellite SAR system specifications, after Raney [97]

However, two operating modes are possible: stripmap mode and wave mode. The entries in Table 1 refer to the stripmap mode and are a fairly complete summary of that mode. More details are available from the ERS web site [1]. For instance the incidence angle varies from 20^o in the near swath to 26^o in the far swath. Also, while the azimuth resolution is stated to be 28 m, this is dependent on the number of looks the data is processed to. According to the web page, the azimuth resolution actually lies between 6 m and 30 m. The special wave mode of operation provides the two-dimensional spectra of ocean surface waves. While such ancillary information can be useful for ship detection purposes, a description of the wave mode is beyond the scope of this report.

A comprehensive description of the Radarsat-1 SAR is given by Raney *et al.* in [98]. They note that the driving design criterion was a standard imaging mode which could be

---

[1]http://earth.esa.int/ers

Mode	Resolution Range x Azimuth (m)	Looks	Width (km)	Incidence Angle (degrees)
Standard	25 x 28	4	100	20–49
Wide(1)	48–30 x 28	4	165	20–31
Wide(2)	35–25 x 28	4	150	31–39
Fine	11–9 x 9	1	45	37–48
ScanSAR (N)	50 x 50	2–4	305	20–40
ScanSAR (W)	100 x 100	4–8	510	20–49
Extended (H)	22–19 x 28	4	75	50–60
Extended (L)	63–28 x 28	4	170	10–23

Table 2: Image characteristics of Radarsat-1 imaging modes, after Raney et al. [98]

used at incidence angles ranging from 20^o to 50^o. However, this design criteria allowed for several other imaging modes at little extra cost. In addition to the Standard mode, the Radarsat-1 SAR has Fine Resolution, Wide Swath, Extended and ScanSAR. We reproduce their table of the image characteristics these modes in Table 2.

The Envisat SAR system is called ASAR (Advanced SAR). An indication of its capabilities is given in Table 1. Those details and the following description were obtained from the Envisat web page [2]. Envisat ASAR was designed to ensure continuity with the image and wave modes the ERS-1/2 AMI, however, unlike the AMI it has a variable viewing geometry and variable transmit/receive polarisations. Thus new modes of operation are possible and the ASAR has enhanced capability in terms of coverage, incidence angles, and polarisation. There are four fundamentally different modes: Image Mode (IM), Wide Swath Mode (WS), Global Monitoring Mode (GM) and Alternating Polarisation Mode (AP). The IM mode is a stripmap SAR mode with high spatial resolution (30 m) and is similar to the ERS AMI image mode. (We note here that the 30 m resolution data is 3-look and a SLC product with 9 m range and 6 m azimuth resolution is also possible). Unlike the AMI however, seven different swaths located over incidence angles ranging from 15 to 45^o are possible. The widths of the swaths vary from 60 km to 110 km depending on the incidence angle. Both HH and VV polarisations can be selected. In WS mode a ScanSAR technique is used to provide images of a wider strip (405 km) with medium resolution (150 m) in HH or VV polarisation. The total swath is built from five subswaths. GM mode is the same as the WS mode except for a lower resolution (1 km) and attendant to that a low data rate. Finally, in AP mode a ScanSAR technique is used but with polarisation changing rather than incidence angle. The result is two images of the same scene in different polarisations combination (HH/VV or HH/HV or VV/VH). The same swaths as Image Mode can be used and the resolution is 30 m. However, the radiometric resolution is reduced compared to Image Mode.

Radarsat-2 is due to be launched late in 2005. Nominal specifications are given in Table 1. Our source was Sanden et al., [108]. Sanden *et al.* note that Radarsat-2 will be capable of all imaging modes of Radarsat-1 as well as incorporating new modes and other significant improvements. In particular they mention improvements in spatial resolution (3 m will be possible), polarisation (QuadPol will be available), orbit control and look

---

[2]http://envisat.esa.int

Mode type	Mode	Swath Width (km)	Looks Rng x Az	Resolution Rng x Az (m)	Incidence Angle (degrees)
SSP or SDP	Standard	100	1 x 4	25 x 28	20–49
SSP or SDP	Wide	150	1 x 4	25 x 28	20–45
SSP or SDP	Fine	50	1 x 1	10 x 9	37–49
SSP or SDP	ScanSAR Wide	500	4 x 2	100 x 100	20–49
SSP or SDP	ScanSAR Narrow	300	2 x 2	50 x 50	20–46
HH only	Low Incidence	170	1 x 4	40 x 28	10–23
HH only	High Incidence	70	1 x 4	20 x 28	50–60
QP	Standard QP	25	1 x 4	25 x 28	20–41
QP	Fine QP	25	1 x 1	11 x 9	30–41
SSP	Multilook Fine	50	2 x 2	11 x 9	30–50
SSP	Ultra-fine Wide	20	1 x 1	3 x 3	30–40

Table 3: Characteristics of Radarsat-2 imaging modes.

direction. The improved orbit control will allow better georeferencing, while the proposed selective look direction is expected to reduce revisit times. In total 11 different modes are possible. These are classified into 3 types: Selective Single Polarisation (SSP), Selective Dual Polarisation (SDP) and Quad Polarisation (QP). SSP allows any of the polarisations HH, HV, VH and VV, to be selected, SDP allows either HH and HV or VV and VH, while QP measures all 4 channels. The QP mode is limited to swath widths of 25 km which will impinge on its utility for ship monitoring. Also the Ultra Fine resolution of 3 m is restricted to the SSP mode. A summary of swath widths, swath coverage and resolution are given in Table 3. This information came from the Radarsat-2 web site [3] and the CCRS report, [110].

We also briefly mention the Japanese ALOS PALSAR system, given that it is highly likely that one of the applications of this system will be ship detection. The ALOS PALSAR launch is due late in 2004. It will be L-band and fully polarimetric with incidence angles in the range 8–60°, resolution in the range 7–100 m, a swath width in the range 40–350 km and a noise equivalent σ^0 of less than −25 dB. See the ALOS web page [4] for more details.

Finally, we note that in describing SAR systems the SAR processor and related image preprocessing algorithms must also be borne in mind. Certain aspects of the imagery quality, e.g. system impulse response function, pixel spacing, peak sidelobe ratio and radiometric resolution, are processor dependent. The ERS web site [5] provides a table with such information for ERS imagery. The paper by Raney *et al.*, [98] has a similar table for Radarsat-1. That such considerations are relevant to ship detection is evidenced in the paper by Aresu and Schwartz, [4], which compares the performance of four different Radarsat ScanSAR processors with particular reference to ship RCS and ship detection.

Various image products are available. The differences between these products relate to

[3] http://www.radarsat2.info/index.asp
[4] http://www.eorc.nasda.go.jp/ALOS/
[5] http://earth.esa.int/ers

factors such as: the number of looks, removal of SAR-related distortion errors, calibration, georeferencing, map projection, etc. Radarsat-1 offers various data products in three main categories: RAW, Path-Oriented and Map-Oriented. RAW data consists of the unprocessed radar signals before image formation. Path-Oriented products are SAR data after image formation but before but before georeferencing and map projection. Single look complex imagery is one of the Path-Oriented products. Map-Oriented products provide imagery in map coordinates with "North Up". Likewise ERS-1/2 offers a range of products (and presumably other systems do too). Among the products offered by ERS, see the ERS web site [6] are:

- Raw data (RAW);

- Single Look Complex (SLC);

- Precision Image (PRI);

- Ellipsoid Geocoded Image (GEC).

We quote from the web page on PRI: *The Precision Image is a multi-look (speckle-reduced), ground range image. It is the product applicable for most users interested in remote sensing applications. The product is calibrated and corrected for the SAR antenna pattern and range-spreading loss: radar backscatter can be derived from the product for geophysical modelling, but no correction is applied for terrain-induced radiometric effects. The image is not geocoded, and terrain distortion (foreshortening and layover) has not been removed.* The GEC product is similar but it is precisely located and rectified onto a map projection.

Radarsat imagery (and presumably other SAR imagery) is also sometimes adjusted using look-up tables. The idea is to improve appearance by applying a range dependent gain correction to 1) optimise the radiometric scaling of the main feature of interest (while optimising the output dynamic range) and 2) compensate for changes in σ^0 with changing incidence angles. A look-up table specifically designed for ocean imaging is available.

2.3 SAR imaging of the ocean

Interest in SAR imagery of the ocean really took off with the launch of the satellite SAR system known as Seasat in 1978. Seasat was specifically designed to image the ocean and provided a wealth of interesting data. Consequently, the topic of SAR imagery of the ocean has a long history and much is now known. Our account here will be necessarily brief and we refer the reader to the extensive literature for further information. A very readable overview is given by Robinson in [106] as is that of Elachi in [29] although the latter is more limited in scope. A more recent, but perhaps too brief summary, is given by Mouchot and Garello in [83]. The report of Ulaby *et al.* [125], is more comprehensive but very technical in places. Alpers, [1], provides a good review of the original work.

Water has a high dielectric constant and hence is highly reflective of microwaves. It follows that active microwave sensors like SAR only record details of the surface layer of the ocean. It is well accepted that for incidence angles above 20^o, Bragg resonance is

---

[6]http://earth.esa.int/ers

the predominant ocean surface scattering mechanism. For lower incidence angles, a facet model with specular reflection can be used (see for instance Elachi [29]). In operational SAR systems, incidence angles are generally above 20^o (although in rough seas the local incidence angle may be less than 20^o in places) and so the Bragg scattering model usually applies. In the **Bragg scattering model**, it is assumed that the backscatter is mainly due to the ocean surface spectral components which generate Bragg resonance. The Bragg wavelength λ_s of the ocean waves responsible for Bragg resonance is given by

$$\lambda_s = \frac{n\lambda_r}{2\sin\theta_i} \qquad (7)$$

where n is the order of the resonance (contributions for $n \neq 1$ are negligible), λ_r is radar wavelength, and θ_i is the incidence angle. This formula assumes the ocean waves are travelling in the direction of the radar line of sight. For common SAR radar frequencies and incidence angles, the Bragg wavelength is in the range 1cm to 40cm. Thus Bragg scattering is from the surface ripples also known as **capillary waves** (length ~ 2 cm) and **short gravity waves** (length ~ 50 cm).

It is clear from (7) that as the radar frequency and local incidence angle change so does the Bragg wavelength. Thus radar frequency and incidence angle will affect the backscatter coefficient σ^0. Insight into what these effects are can be deduced from knowledge of the ocean wave spectrum $W(k)$. For instance, it is known that $W(k)$ decreases rapidly with increasing wave number k (estimates range from $k^{-4} tok^{-8}$ in the Bragg region). Since smaller wavelengths correspond to large wave numbers it follows that radar backscatter should decrease rapidly with increasing incidence angle. This is borne out by empirical data. Similarly, radar backscatter decreases with increasing radar frequency. Although the Bragg scattering model can be used in this way to explain many of the properties of σ^0, it has limitations. In particular, it does not explain the effects of radar polarisation. Ocean reflectivity is usually larger for VV polarisation than for HH polarisation with the size of the difference dependent on the radar incidence angle. A very informative figure describing this difference is given by Vachon *et al.* in [127]. For a more detailed account of such topics, see for instance the book by Long [72].

Another key factor in interpreting SAR imagery of the ocean is the near-surface wind field. The surface wind has a direct effect on the surface ripples (capillary waves) and hence its effects appear in SAR imagery. In fact, wind stress on the ocean surface is what generates the surface ripples. Moreover, ripples are short lived and so reflect quite closely the current wind conditions. Thus higher wind speeds cause a rougher ocean surface, increasing the amplitude of the waves responsible for Bragg scattering and hence increasing the radar backscatter. Conversely, lower speeds result in a smoother surface and darker radar backscatter. Roughly, [97], for a given radar and incidence angle there is a first order power law dependence $\sigma^0 \sim |u_s|^\gamma$ where u_s is the wind speed and the exponent γ is a function of radar wavelength and other such parameters. Interestingly, this dependence of σ^0 on the local wind speed is strong enough that it can be used as a means of estimating wind speed. All that is needed is an accurate model relating σ^0 to wind speed and direction and the other relevant factors.

As an example of the type of knowledge available on the ocean backscattering coefficient, we cite the empirical model used by Vachon *et al.* in [127] for ship detection performance assessment. It predicts that

13

- σ^0 increases with increasing wind speed

- σ^0 is largest when the wind is blowing towards the radar

- σ^0 is smallest when the wind is blowing across the radar look direction

- σ^0 for C-band VV is larger than σ^0 for C-band HH for all wind speeds and directions

- σ^0 decreases with increasing incidence angle

- σ^0 for C-band HH decreases more rapidly with increasing incidence angle than σ^0 for C-band VV

See also the model used by Vachon *et al.* in [129] for extracting wind speed from σ^0 values.

The ocean swell and other large scale wave structures such as ship wakes are also visible in radar imagery. This is because they modulate the Bragg wave structures. The modulation is usually described in terms of the **two-scale model** (also known as the composite surface model). In this model, the ocean surface is seen as being composed of facets (or large scale waves) with the fine scale Bragg wave structure superimposed. The Bragg wave structure determines the surface roughness and hence its radar reflectivity while the facet (or large scale wave) orientations determine the local angle of incidence and hence determine which Bragg frequencies are relevant. This model it is most applicable for low sea state conditions and mid-range angles of incidence, but it is widely used and suffices for our purposes.

It is generally agreed that there are three ways in which ocean waves modulate the Bragg wave structure:

- Tilt modulation - tilting of the ocean surface by the larger scale waves,

- Hydrodynamic modulation - interactions between the larger and small scale waves,

- Motion effects - ocean surface motion due to the passage of larger scale waves.

Each of these mechanisms will affect the amplitude of the Bragg waves and hence will have a direct impact on the ocean RCS. The first two effects are fairly straightforward to understand. The third is more difficult and is comprised of two effects: velocity bunching and azimuth image smearing.

Tilt modulation is the name given to the tilting of the ocean surface by the larger scale waves. It is the purely geometric effect in which the shape of large waves determines the local incidence angle at the patch of ocean surface being imaged. Changes in local incidence angle lead to changes to the relevant Bragg frequencies and hence to changes in the backscatter intensity. Clearly then, the effects of tilt modulation are in phase with the large scale waves. **Hydrodynamic modulation** is the name given to the hydrodynamic effects which result in the non-uniform distribution of ripples on the larger waves. Robinson, [106], describes a possible mechanism for hydrodynamic modulation. In his explanation, the large scale wave motion creates surface currents which interact with the ripples. The currents converge at wave crests, increasing the ripple amplitudes, and diverge at wave troughs, decreasing the ripple amplitudes. Again, this effect is in phase with

the large scale waves. **Motion effects** are caused by the motion of the ocean surface. Here we are not referring to the evolving dynamic nature of the surface which results in a lack of scene coherence but rather to the predictable orbital motion of the surface due to the action of waves. As discussed in the last section, such motion can result in azimuth image shift and smearing. Fortunately, the ocean moves slowly enough that these effects are not too severe and SAR ocean imagery shows much of the structure present. In fact, as we explain next, motion effects can aid the imaging of ocean waves through a process known as velocity bunching.

According to the two-scale model, the ocean surface moves at two scales — ripple waves and ocean swell waves. Motion of the surface ripples Robinson, [106], states that the first of these motions is too small to matter. However, the ocean swell wave motion is significant. The surface motion associated with ocean swell waves is orbital. Thus parts of the wave rise towards the SAR sensor, other parts fall away and yet other parts have no velocity component in the direction of the sensor at all. This results in varying azimuth shifts throughout the wave cycle and the combined effect is known as **velocity bunching**. For range travelling swell waves, the crests and troughs lie parallel to the azimuth direction and so velocity bunching has no effect. Any other waves will have an azimuth travelling component and velocity bunching comes into play. Its effect is to bunch the scatterer facets over the wave crest making it brighter and spread them thinly over the wave trough making it darker or visa-versa depending on the direction of the wave travel.

Other effects also need to be taken into account when analysing SAR imagery of the ocean surface and swell wave patterns in particular. We have already mentioned the evolving dynamic nature of the surface which results in a lack of scene coherence and a consequent degrading of azimuth resolution. Vachon and Krogstad, [130], note that this effect is responsible for the so called **azimuth cutoff** which is characteristic of SAR ocean wave image spectra. The term azimuth cutoff refers to the inability of SAR systems to properly image azimuthal waves with short wavelengths. Vachon and Krogstad note that the degree of velocity bunching nonlinearity and the azimuth cutoff are proportional to the SAR range to velocity ratio as well as depending on the actual wind and waves conditions. The ERS-1 SAR system with $R/V \approx 115s$ has trouble imaging azimuthal waves whose wavelength is less than about 200 metres. There are also noncoherent imaging effects which come into play, for instance: scanning distortion and look misregistration. These effects are a consequence of the translation of the ocean wave pattern at the wave phase velocity. A discussion of these effects may be found in, [130], for instance.

Much research effort has been invested in the problem of deducing the directional ocean wave directional spectrum from SAR observations of the wave field. The theoretical developments have a long and complicated history but the details are still not fully resolved, see for instance [130]. It is clearly beyond the scope of this report to try and unravel the research on this topic. However, as a representative recent example, we cite the work by Dowd *et al.*, [27]. As far as ship detection is concerned, extracting wave (and wind) information from SAR imagery is potentially important because: it helps determine how well the detection algorithms will do; it helps in the selection of the algorithm parameters; it effects ship routing and hence provides higher level detection information - certain boats will not be out in certain sea states; etc. However, none of the ship detection systems in the literature have reached this level of sophistication yet.

Waves are not the only oceanographic phenomena visible in SAR imagery. Any process which affects the surface Bragg resonant wave structure can be imaged with SAR. There are three basic mechanisms by which this occurs: through direct contact with the surface (surface processes), through roughening of the surface and wave generation by near surface winds (atmospheric processes), and through hydrodynamic interaction with near surface currents (sub-surface processes). Such processes will interact and their combined effect on the Bragg wave structure will be determined by a mixture of their effects on near surface winds, near surface currents and surface slicks. It follows that the resulting SAR image will show a mixture of oceanographic processes and unravelling them may be difficult.

The following is a comprehensive (although likely not complete) list of processes and objects which may be seen in SAR imagery of the ocean.

- Atmospheric winds and storm fronts

- Heavy rain

- Gravity wave fields (arising from the local wind field)

- Ocean swell (arising from distant wind fields)

- Ocean currents and eddies

- Overfalls (breaking waves caused by conflicting winds and currents)

- Internal ocean waves

- Ocean bottom topography

- Man-made and natural slicks and films

- Sea ice

- Ships and Ship wakes

- Man-made and natural objects such as shoals, islands, oil rigs, seaweed, etc.

Many of these processes and objects will be of relevance to ship detection algorithms since they can cause false alarms. We cite some examples from the literature next.

Many authors report that the variations in ocean backscatter caused by variations in surface wind speed and direction can cause ship detection problems. For instance, both Kourti et al., [61], and Schwartz et al., [114], report that local wind effects can be mistaken for ships. Similarly, Lombardo and Sciotti, [70], mention the transitions regions between different wind conditions and low wind spiral marks as being problems. They also add that in low wind speed regions, the backscatter is high variability and the mean intensity is low which results in spots of bright ship-like backscatter. But there are many other processes and objects which cause problems too. Vachon et al., [134], report that oceanographic phenomena (such as breaking waves), outlying rocks, shoals and islands can cause false alarms in the OMW. Yeremy et al., [142], report that small islands, buoys and ocean based oil rigs also cause false alarms in the OMW but add that islands (with the exception of very small islands) and oil rigs can often be distinguished from ships visually. Likewise,

Gower and Skey, [43], report that rocks and small islands appear as bright ship-like point targets. Rey *et al.*, [102], mention oceanic/atmospheric features such as fronts and storm cells. Ouchi *et al.*, [93], mention current boundaries and breaking waves. Lombardo and Sciotti, [70], report that inhomogeneous ocean areas cause problems and add land area effects and bathymetry effects to the wind effects already discussed. Cusano *et al.*, [25], report that surface wind speed and direction, sea currents and coastal effects can cause high variability of the backscattered intensity which causes problems for ship detectors. They add that this variability is enhanced by other factors such oil spills , pollution, and high densities of ships and wakes. Ringrose and Harris, [104], mention that low wind speed causes dark patches and when mixed with normal wind speed regions the resulting variation in backscatter causes problems for conventional CFAR detectors. Eldhuset, [31], reports that eddies, fronts and internal waves can be mistakenly identified as ships and ship wakes. Wahl *et al.*, [138], report that some of the ocean features in ERS-1 imagery with similar signatures to ship wakes are: ocean current shears; upwelling and internal waves. They also report that ocean wave patterns can cause strong backscatter which may be mistaken for ship backscatter.

2.4   SAR imaging of ships and ship detectability

We begin this section with a general discussion of SAR imaging of ships and the various factors involved. Because of the diversity and complexity of such factors not much work has been undertaken on modelling ship backscatter for detection purposes. In fact only one such model has been used in the literature. We then turn to models which have been used for estimating probabilities of detection. This leads to a discussion of detection performance and ship detectability. We complete the section with a brief discussion of the various factors affecting ship detectability. The latter is a topic in its own right and beyond the scope of this report. We note that previous work on this topic has been undertaken at DSTO, see for instance the reports by Kruzins and van Antwerpen, [2] and [3], and the related conference papers, [28] and [62]. We also note that there is an extensive literature on fisheries monitoring which is more concerned with ship detectability than detection algorithm development.

Ships are typically constructed from large flat metal sheets and hence are usually radar bright and therefore detectable in SAR imagery. While backscatter from the ship itself is the usual basis for ship detection algorithms other factors can also be taken into account. Ships often have a wake and further, the wake can be more noticeable than the ship itself. We discuss wakes in the next section. A radar shadow is possible too. However no mention of this is made in the ship detection literature and we assume that radar shadows are not usually present. (Obviously this is true for low resolution imagery where the image resolution cells are larger than the ships.) We also remind the reader here that any radar shadow will be at the head of the wake and offset from the ship. Higher level information may be relevant too. For instance: ocean going ships will tend to be larger and made of metal whereas coastal boats will be smaller and may be made from wood or fibreglass; oil slicks are often associated with ships; and the location of shipping channels and regular routes taken by ferries may be useful.However, none of the automatic detection systems have reached this level of sophistication yet.

Askari and Zerr, [7], provide a brief but useful summary of factors affecting the SAR imaging of ships. In their summary, they cite a report by Asknes which could not be obtained. They state that backscatter from ships is determined by several scattering mechanisms including: direct reflection from areas perpendicular to the radar beam, corner reflections and multiple reflections from the ship and sea surface. They go on to note that other factors which come into play are: construction material, superstructure, incidence and aspect angles, radar frequency and polarisation. As an example, they report that studies using multiple frequencies have shown ship backscatter is enhanced by higher frequency SARs. Another example is Asknes' finding that ships travelling in the azimuth direction have larger RCSs. Askari and Zerr also note that motion effects (speed, roll, heave) are important. These effects reduce the coherence of the backscatter resulting in smearing and reduction in peak pixel intensities. Yeremy et al., [142], make similar general comments about the SAR imaging of ships and cites several studies on ship RCS. They add spectral leakage and sidelobes to the list of complicating factors. Kourti et al., [61], provide empirical data on how RCS varies under different imaging conditions. The data verifies that a vessel's heading is important in this respect and RCS is lower for a range travelling vessel than for an azimuth travelling vessel. They warn however that it is difficult to draw such conclusions due to the large number of variables involved. Whether the vessel is laden or unladen may be relevant too. For instance, a fully laden vessel will sit lower in the water and hence will tend to have a smaller RCS. Note that our use of the term ship RCS in the above is slightly misleading. The relevant quantity for SAR imagery is the incoherent average value of σ^0 over the ship pixels whereas total ship RCS is the coherent average.

Not a lot of modelling of ship backscatter has been undertaken in the ship detection literature — probably because too many factors are involved making it too difficult. Mostly people do not take the ship itself into account and simply search for anomalies in the background. In fact the only ship detection algorithm found in the literature which is truly based on a ship backscatter model is the generalised likelihood ratio (GLRT) detector used by Sciotti, Pastina and Lombardo, [117] and [116]. While their work is mainly focused on polarimetric SAR imagery, their work arose from earlier work by Brizi et al., [10], on single channel imagery. The model used by Brizi et al. is very simple and consequently probably not very realistic. A review of this model is given in Section 5.6 and its extension to polarimetric data is discussed in Section 40. We merely comment here that the imagery is SLC and the ship pixels are assumed to be samples from a zero mean complex Gaussian random process which is independent from pixel to pixel but correlated (with restrictions) within the polarimetric channels. This type of model is a commonly used for clutter samples from SLC imagery but its application to target samples is questionable. Actually, Brizi et al. initially propose that the ship pixels are not zero mean which makes more sense but drop this refinement in order to calculate the likelihood ratio. Also, while the independence of samples from pixel to pixel is reasonable for low resolution imagery it is probably questionable for high resolution imagery. This simple circular complex Gaussian target model has also been proposed by Lopes et al., [73], in the context of a likelihood ratio detector.

While it is rare for ship target models to be used for detector design as such, they are often used for theoretical detector performance assessment and for estimating probabilities of detection in particular. This feeds back into detector design in that knowledge of how

the probability of detection varies with the detector threshold allows the threshold to be set "correctly". Before we discuss the detectors in which this is done we mention a third use of ship target models. Gower and Skey, [43], use a simple Gaussian "blob" model for the collection of ship target pixels in order to extract discrimination/classification features. The model is parametric and is fitted to the SAR data at points of interest. We discuss it in the section on discrimination algorithms.

Lombardo and Sciotti, [70], assume ship backscatter is contained entirely in one resolution cell and model it with a chi-squared distribution with order parameter ν_T for single look imagery. This becomes the Gamma distribution with order parameter $L \times \nu_T$ for L-look imagery. They note that the special cases of Swerling I and III target models correspond to the order parameters $L \times \nu_T = 1$ and $L \times \nu_T = 2$. Lombardo and Sciotti point out that in high resolution imagery single pixel target models are no longer valid and it is better to use all the target pixels. They add however that such schemes require knowledge of target size, shape and orientation or the use of parametric matched filters and are computationally expensive. For this reason they retain the pixel by pixel approach even in high resolution imagery. The model is not used for detector design as such. Rather, it is used for the theoretical analysis of probabilities of detection. Lombardo and Sciotti note that in addition to threshold setting, their analysis also allows the comparison of SAR imagery of different spatial and radiometric resolutions. Likewise, Ferrara *et al.*, [33] and [34], use a Rician pdf to model the distribution of ship target pixel intensities in the case of 1-look processing and a gamma distribution in the case of multi-look data. Again, the models are not used for detector design but rather for the theoretical estimation of probabilities of detection and consequently setting the detector threshold.

A model which has been discussed extensively is that used by Vachon *et al.* in [127] (see also [14]). This model is used for theoretically estimating the performance of ship detection algorithms and thereby assessing ship detectability. The model is very simple. It is applied to low resolution imagery and assumes ship backscatter is contained within one image pixel. It relates ship RCS to its weight in tons and by further analysis to its length. Thus the ship RCS σ is assumed to satisfy

$$\sigma = D = 0.08l^{7/3} \tag{8}$$

where D is the ship displacement in tons and l is its length. The assumption that $\sigma = D$ was offered by Skolnik in the absence of any better information. The relation $D = 0.08l^{7/3}$ was obtained empirically by regression on real ship data. Clearly, the model (8) takes no account of fundamental variables like radar incidence angle, ship orientation, construction materials, superstructure, etc. and so has limited application. Such shortcomings are acknowledged by Vachon *et al.*. They add that the model has the correct order of magnitude but tends to underestimate RCS, especially with increasing incidence angle. Yeremy *et al.*, [142], also comment on the limitations of (8) and cite several relevant references. However, as Rey *et al.* point out in [103] the model is easy to apply and as explained next, Vachon *et al.* do not use it for detection algorithm design but rather, for theoretical performance assessment.

The formula (8) is part of a statistical model which Vachon *et al.*, [127], use to predict ship detectability. Their aim was to facilitate comparison of the various Radarsat beams and modes for ship detection. In addition to (8), the full model also relies on assumptions

about the SAR image probability density function (pdf) and the relationship between ocean backscattering coefficient σ^0, wind speed and SAR viewing geometry. The SAR image pdf used is the multi-look K-distribution. We discuss it below in the section on the OMW ship detector. The backscattering coefficient σ^0 is needed to scale the image pdf. Vachon *et al.* note the lack of good physics based models for σ^0 and use instead an empirical model. This model was developed for ERS-1 data but is adapted to Radarsat data. When all the ingredients are combined, the model allows for estimation of the minimum detectable ship length for various radar parameters and ocean conditions. Vachon *et al.* verify many of the model's assumptions and predictions on real data. The model predicts that for Radarsat SAR imagery:

- Beam mode Standard 1 has better ship detection performance than ERS-1 due to the decreased clutter level of HH polarisation,

- Ship detection performance improves with increasing incidence angle due to the reduction in clutter level while ship RCS remains constant,

- Ship detection performance is best for fine beam modes due to their large incidence angle and fine resolution,

- For ScanSAR modes, the detection performance is best for large incidence angles, but worse than the standard beam modes due to the coarser resolution,

- For large incidence angles, the ship detection problem becomes the detection of bright point targets against a noise background.

Vachon *et al.* conclude that Radarsat ScanSAR Narrow Far with its 300km swath, is a good compromise between ship detectability and swath coverage. They also note that smaller ships may be detected with decreasing wind speed, increasing incidence angle and finer resolution. In a later paper based on this work, Vachon *et al.*, [134], recommend ScanSAR Narrow Far and perhaps ScanSAR Wide for surveillance purposes, and beams Wide2, Wide3, Standard4–7, Fine1–5 and ExtendedHigh1–6 for tracking purposes.

A modified version of the model which takes account of the radar incidence angle has been used by Askari and Zerr, [7], for a similar analysis of Radarsat data. It has the form

$$\sigma = 0.08 R(\theta) l^{7/3}, \qquad R(\theta) = 0.78 + 0.11\theta \qquad (9)$$

where θ is the radar incidence angle. Their findings are similar to those of Vachon *et al.*. The modified model has also been used by Olsen and Wahl, [90] and [89], for an analysis of the Envisat ASAR sensor. Their findings are also similar. However, based on preliminary results, they recommend the Envisat Alternating Polarisation mode with VV and VH channels. They note that at steep incidence angles vessels should be easier to detect in the VH channel since the clutter levels are lower than in the co-polarised channels while the ship RCS remains the same.

Much of the early work on ship detection and detectability was applied to Seasat and ERS-1 SAR data. We take the paper by Wahl *et al.*, [138], as representative of the main findings of this work. Wahl *et al.* analyse ERS-1 SAR Fast Delivery data with a resolution of about 30 m. Their main findings are that: ships longer than 120 m are generally seen in

the imagery under most weather conditions; and ships smaller than 100 m are sometimes lost because of strong ocean backscatter in high winds (greater than 10 m/s). They note that ERS-1 imagery is sensitive to wind speed. They present a table of wind speed versus the lower length of visible ships. It shows that: most ships are visible when there is no wind; ships longer than 50 m are visible in wind speeds of 4 m/s; and ships longer than 100 m are visible in wind speeds of 10 m/s. They also observe that the far range is better for ship detection than the near range because ocean clutter is lower. Small ships may become invisible in the near swath and this effect has clearly been observed with fishing vessels. They also note that detection thresholds need to be increased when wave patterns are present. Wahl *et al.* also include many comments on wake types and sizes. They note that wakes are more susceptible to high wind speeds than ships. They also note that wakes without ships are rare but the opposite is not the case, so the ship itself is the more robust detection feature. They conclude the paper with the comment that the main limitation to ship detection is probably not the algorithms but the strong ocean backscatter which may hide the ship backscatter.

Both the work of Vachon *et al.* and Wahl *et al.* discussed above have assessed ship detectability in terms of ship length. We have already noted that this is a simplification of the real situation and that many other factors are likely to come into play. The paper by Friedman *et al.*, [38], provides concrete evidence of this with empirical data on ship RCS and ship detectability. The work was aimed at validating a CFAR vessel detection algorithm and consequently well ground-truthed data was used. The results clearly show that ship length alone is not enough to predict ship detectability. Friedman *et al.* suggest that other important factors are ship orientation and local wind and wave conditions.

Many authors have reported the importance of incidence angle and wind speed. Ocean clutter levels are higher at small incidence angles and high wind speeds. Thus, Rey *et al.*, [103], point out that the far swath of Radarsat is good for ship detection while the near swath is good for wake detection. Kourti *et al.*, [61], add that when small incidence angles are combined with high wind speed, detectability problems are severe. On the other hand Robertson *et al.*, [105], find that detectability improves (modestly) with increasing incidence angle but that there is little correlation with wind speed. Robertson *et al.* add that for Radarsat ScanSAR Narrow Far data, ships down to a size of 30–40 m could be seen visually in moderate wind speeds of 5–10 m/s. Gower and Skey, [43], report that larger incidence angles give better ship discrimination in high wind conditions. They note that SAR image intensity depends on the direction of the wind relative to the SAR look direction and add that the presence of wind streaks can indicate wind direction.

The importance of radar polarisation and frequency have been noted too. For instance, in [109], van der Sanden *et al.* assess the applications potential of Radarsat-2. They comment that HV polarisation of Radarsat-2 should provide better ship detectability in the near range (i.e. small incidence angles) than the HH or VV polarisations. This is because the backscatter from the ocean should be smaller for HV than HH or VV and consequently ship/sea contrast should be higher. Similarly, Dinguirard *et al.*, [26], note that the JERS-1 (L-band HH 35°) has a better ship detection capability than ERS-1, and Yu and Litovchenko note that ALMAZ-1 (S-band HH 20–60°) images ships better than ERS-1. They claim that ships of 10–15 m length are visible under wind speeds of 10–15 m/s as well as in calm wind conditions of 0–3 m/s. Eldhuset, [32], notes that the short wavelength and small incidence angles of the ERS SAR system makes it strongly

sensitive to ocean backscatter in high wind conditions and this limits its potential for ship detection.

Ship construction material has also been mentioned too. In particular, it was observed by Kourti *et al.* in the paper [61] that small fibreglass and wooden fishing vessels were not discernible in Radarsat imagery even in calm sea conditions because of their low dielectric characteristics. They comment further that steel structures on such vessels may be visible but their reflectivity is aspect angle dependent and therefore unreliable.

2.5   SAR imaging of ship wakes

Ship wakes are also detectable in radar imagery. They have several components which modulate the ocean backscatter coefficient. These wake components may be either brighter or darker than the surrounding ocean, depending on the radar imaging parameters and the sea state. As mentioned in the introduction, we are concentrating on ship detection algorithms in this report. It is anticipated that the detection of ship wakes will be the topic of a companion report. However, the question still arises: Should one detect the ship first or the wake first? The advantages of detecting ships first include:

- Stationary and slow moving ships have no wake.

- Wakes without ships are rare in ERS-1 imagery, [138] and [32].

- Ship backscatter is robust and independent of sea state, [32].

- Wakes are very variable and depend on the sea state and ship speed and direction;

- Wakes are often not visible and especially at large incidence angles and long wave-lengths. For example, JERS-1 does not show wakes [32].

- Wakes of maneuvering ships have variable shapes.

- Ship detection with a CFAR algorithm is fast and simple whereas wake detection with the Radon transform is slow.

- Good detection results have been reported for ship backscatter, [134].

For these reasons, most systems in the recent literature detect the ship first. However, there are advantages to detecting ship wakes first

- The wake (when present) can be more noticeable than ship itself since wakes are much larger and can be several kilometres long.

- Azimuth motion shift of the ship may place it in an area of high clutter making detection harder.

- Wakes indicate the ship's true locations (no motion shift).

Many of the properties of ships and wakes just listed are dependent on the radar parameters and ocean conditions and so the decision of which to detect first is likewise dependent on the radar parameters and ocean conditions. A good discussion of such issues is given by Eldhuset in the introduction of [32].

A good example of how radar parameters affect the imaging of ship and ship wakes is provided by the comparison of Radarsat data (C band and HH polarisation) with ERS data (C band and VV polarisation). Wackerman *et al.* [135] note that most early work on spaceborne SAR for ship detection focused on ERS-1 data. The ERS-1 system was VV and used small incidence angles and hence ship wakes showed up well. In fact, wakes were larger and more distinct than the ship return itself and so detection systems were designed to find wakes. However, the Radarsat system was HH and used large incidence angles. HH ocean clutter is smaller than VV ocean clutter and decreases more rapidly with incidence angle and for this reason ships show up well in Radarsat imagery and ship wakes are rarely seen. Consequently, Radarsat data is more suitable for detecting the ship target itself.

It has recently been found that VH polarisation provides the best ship-sea contrast (at low incidence angles). Consequently, the recommendation of Sanden *et al.*, see [110] and [109], for ship surveillance with Radarsat-2 is a dual polarised ScanSAR wide mode data set composed of VV and VH channels (fully polarimetric data will not be available in ScanSAR wide mode). The VH channel should provide bright point targets against a dark background while the VV channel should provide sufficient ocean clutter for wake analysis. Olsen and Wahl, [90], likewise recommend the use of Envisat's Alternating Polarisation mode with VV and VH channels although they do not specifically mention wake detection in the VV channel.

3   Land masking

Most ship detection systems include an initial stage in which the land areas are masked out of the image being processed. Detailed accounts are in [135], [32], [34] and [107], and passing mention is made in [61], [47], [48], [25], [68] and [105]. Land masking is important not only for the obvious reason that only ships in the ocean are of interest but also because ship detectors can produce high numbers of false alarms when applied to land areas. Dealing with these false alarms can overtax detection systems. For similar reasons, it may be necessary to mask out other image areas. For example, Henschel *et al.*, [47], note the need to mask out beam seam and nadir ambiguities in Radarsat imagery and Rey *et al.*, [102], note the need to remove "dark stripes" of constant valued pixels. Another example is the presence of ice sheets in the polar regions. For this reason, the OMW accepts general user defined masks as well as land masks, see [111].

A common approach to land masking is to register the SAR image with existing geographic maps. While this is the easiest approach it is not perfect. In the case of satellite imagery, the orbital parameters are not known precisely and registration errors occur. Tidal variations are obviously also possible and rocks and small islands may not be marked on the geographic maps. Thus some land areas will be marked as ocean and vice-versa.

For Radarsat imagery, Wackerman *et al.*, [135], note that registration errors are typically of the order of hundreds of metres but in the worst case can be up to 2 kilometres. For this reason they use a ±2 kilometre buffer zone around the coastline. In order not to miss any ship detections they apply their detection algorithms inside the buffer zone. They deal with the higher false alarm rate in this zone by including additional constraints. A similar 2 kilometre buffer zone was used by Robertson *et al.* [105]. Kourti *et al.*, [61], note consistent geolocation errors of about 1 nautical mile for Radarsat imagery.

In [32], [30] and [138], Eldhuset and Wahl *et al.* also use geo-registration to locate coastlines in ERS-1 SAR imagery. The initial step is to calculate geodetic longitude and latitude for the imagery from satellite orbital parameters and other quantities, [30]. As with Radarsat data, this process has inaccuracies in the order of hundreds of metres. In order to correct these inaccuracies, Eldhuset and Wahl *et al.*, [32] and [138], use a correlation technique. A SAR image is simulated from digital terrain models and then overlayed on the real SAR image. A correlation is then performed using the FFT. This improves the registration accuracy to better than 100 metres. A subtle but important point is that in mountainous areas the "lay-over" effect needs to be taken into account when simulating the SAR image.

Some systems allow user intervention in the geo-registration process. Eldhuset, [32], notes that sometimes the correlation of the map with the SAR image does not work in which case the user can perform the correlation manually. Similarly, Rye *et al.*, [107], have a manual mode of operation. They note that in this mode, the registration is better and a smaller buffer zone may be used around the coastline.

Another approach to land masking is to use algorithms which automatically detect coastlines. Among the reasons for considering such an approach are: the geo-registration errors just mentioned; satellite orbital data may not be available; inaccuracies in existing maps; shifting coastlines; and not all small islands and rocks are mapped. Moreover, if the littoral zone is of particular interest then tidal variations in the coastline will need to be taken into account. We add here that many rocks and shoals may only be exposed at low tide.

Only one ship detection system mentions the automatic detection of coastlines. That system is the Alenia Areospazio ship detection system described by Ferrara and Torre in [35] and [34]. Their algorithm is based on the work of Lee [64]. The basic steps involved are: filtering to remove speckle, application of an edge operator, edge map dilation with a mean filter, thresholding of the edge map histogram, and finally applying a contour following algorithm. They include a minimum land area constraint to prevent the ships themselves being masked out. Ferrara and Torre apply the technique to SLC SIR-C L-Band SAR imagery processed to a range of different number of looks and resolution [34]. They state that there were no problems with detecting coastlines correctly.

While only one ship detection system uses automatic coastline detection, we note that there is an extensive literature on the topic. It is beyond the scope of this report to delve into that literature but we do mention the following. Van Der Sanden *et al.*, [109], recommend the use of VH polarisation for coastline extraction. The paper [65] by Lesage and Gagnon proposes a level set based active contour approach. The report by Yeremy *et al.*, [144], considers fully polarimetric data and lists the following 4 basic steps: 1) apply desired contrast enhancement and filter to remove speckle; 2) interpolate or subsample for

better resolution; 3) use clutter statistics to extract shore-line; 4) use morphologic filters to remove extraneous shore-lines.

4 Preprocessing

The aim of preprocessing is to make the subsequent detection stages easier and hence more accurate. Consequently, the type of preprocessing used will depend on the nature of the subsequent detection stages and preprocessing and detection should be considered in unison. For instance, while speckle filtering enhances the visual quality of SAR imagery, there are no guarantees it will improve the performance of later detection algorithms unless they are specifically designed for speckle filtered imagery. Further to this point, we note that the commonly used CFAR detectors are based on assumptions about the background statistics which do not include the changes speckle filtering brings about and therefore they are not suitable (in theory) for speckle filtered imagery.

In their earlier work on ship detection, [34] and [35], Ferrara *et al.*, recommend the use of speckle filters as a preprocessing step. In [35], they cite in particular the paper [74] where a speckle filter which is not only adaptive but also takes account of image structure is described. In [34], they merely state, without going into details, that a range of user selectable adaptive speckle filters is available. Following speckle filtering, ship pixels are detected using a threshold whose value depends on the background values. In later work, [33], Ferrara *et al.* have discarded speckle filtering as a preprocessing step.

Eldhuset [30] also used speckle filtering as a preprocessing step. He reports that the most robust results were obtained with Frost's adaptive Wiener filter [40]. He provides full details of the filter. Eldhuset follows the speckle filter with a Butterworth highpass filter to enhance the targets and then applies a threshold to detect them. In later work Eldhuset [32] abandons this approach and instead uses a CFAR detector without any preprocessing. Interestingly, Eldhuset implies that the earlier algorithm was more accurate but was dropped because it was too computationally expensive for his application.

Several other papers mention speckle filtering, [139], [8] and [5] but the subsequent detection methods are not particularly convincing.

The system reported on by Rye *et al.* [107] uses a preprocessing step which normalises the imagery by dividing each pixel value by its local mean as calculated over a large window of 30x30 pixels. The effect of this is to remove large image structures while retaining small (point or linear) features. Rye *et al.* then detect potential ship pixels by applying a simple threshold. In fact, as will become clear in the next section, this combination of normalisation and threshold detection is effectively just a cell-averaging CFAR detector. However, Rye *et al.* point out that the normalisation can be followed by any detection algorithm which assumes the background is homogeneous.

Calibration of the imagery to units of NRCS or some equivalent is often necessary in order that absolute thresholds may be used or physical quantities of interest estimated. For instance, this is the case in [135] where a minimum energy constraint on the ship return is used to eliminate false alarms due to noise spikes and in [127] where the ocean RCS is used in estimating the detectability of ships on the basis of a ship scale parameter.

Ferrara *et al.*, [33] preprocess SLC SIR-C/X-SAR imagery by applying spatially variant apodization (SVA). SVA is a non-linear technique for suppressing SAR processor sidelobes with the benefit of not reducing the image resolution or altering the phase information. Ferrara *et al.* report that this step reduces the false alarm rate of their CFAR detectors significantly.

Lopes *et al.* [73] recommend the use of a spatial whitening filter (SWF). They note that when clutter samples are correlated, the arithmetic mean of intensity is not a minimum-variance estimator of power levels and hence they suggest the use of SWF. The SWF was introduced by Larson *et al.* in [63]. Larson *et al.* show that in general the SWF is a superior method of speckle reduction compared to noncoherent spatial averaging. However, they also note that if the SAR data is spatially white then the SWF is identical to noncoherent averaging.

5 Prescreening — single channel imagery

Algorithms which search the whole image for potential ship pixels are referred to as *prescreening algorithms*. Experience has shown that no one algorithm or even a combination of algorithms will do this perfectly — some ship pixels will be missed and some false alarms will be generated. Thus a follow up stage is often included in target detection systems in which the false alarms are identified and removed. This follow-up stage is referred to as the *discrimination stage*. We discuss discrimination algorithms in a later section. There are two important factors which exacerbate the need for a discrimination stage. Firstly, prescreening algorithms often have adjustable design parameters which control a tradeoff between the missed detections and the false alarms. Such parameters are usually set to achieve a high probability of target detection but there is an attendant high false alarm rate. Secondly, near real time processing is usually a requirement of operational systems and so speed is often a consideration in designing prescreening algorithms. As a consequence of this prescreening algorithms are often not as accurate as they may otherwise be. Speed considerations are not such an important factor in designing discrimination algorithms since they only process the output of the prescreening algorithms and not the whole image.

In this section, we review prescreening algorithms with application to single channel SAR imagery — multi-channel imagery is considered in later sections. In general, ships show up in single channel SAR imagery as individual pixels or groups of pixels which are unusually bright compared to their surroundings. Not surprisingly then, the prescreening algorithms of most ship detection systems for single channel SAR imagery involve a method of searching for bright anomalies in the ocean clutter. While pixel brightness is a distinguishing feature of ships, other features such as texture may also be taken into account. As we shall see though, most ship detection algorithms are based on modelling the background statistically and then finding individual pixels or small groups of neighbouring pixels whose brightness values are statistically unusual. Modelling ship targets and taking their statistics into account is generally too difficult owing to the large number of factors which come into play.

In designing and implementing prescreening algorithms, care needs to be taken with

the type of SAR imagery to which they are applied. Algorithms designed for one type of imagery may not be appropriate or may need to be adjusted for a different type of imagery. Resolution is obviously a critical factor in this respect, but other factors, such as whether the imagery is complex, intensity or amplitude and the number of looks are important too.

5.1 Global threshold algorithms

The simplest way to search for radar bright objects is to fix a global threshold and then declare any pixel intensities which are above this threshold as targets of interest. This is the approach used by Lin *et al.*, [68] and [67], for ship detection in ERS SAR PRI images. In order that other types of imagery may also be used, Lin *et al.* first calibrate the imagery so that the pixel values are normalised radar cross section (i.e. RCS per unit area). A simple fixed threshold is then applied. Lin *et al.* provide no details of how this threshold is set. As with other prescreening algorithms, not all pixels detected by this process will be true ship pixels and further processing is required to remove such pixels.

The RCS of a ship is dependent on many factors including the construction materials, the viewing aspect angle and the radar resolution. Thus a large metal ship lying broadside to the radar will have a large RCS while a small wooden boat that is end on will have a much smaller RCS. In low resolution surveillance imagery, a ship may only occupy part of a resolution cell (pixel) and consequently the RCS of the resolution cell will be lower than expected. For this reason, global thresholds are rarely used and it is much more common to apply an adaptive threshold. Adaptive thresholds are designed to look for pixels which are bright, but only relative to their surroundings.

5.2 Adaptive threshold algorithms

Adaptive threshold algorithms are the most common prescreening algorithms for target detection in radar imagery. In this section, we provide a general introduction to the two most commonly used adaptive threshold algorithms for ATD in SAR, namely, the two parameter CFAR detector and the cell-averaging CFAR detector. In the next section, we continue with a discussion of how these and other adaptive threshold algorithms have been used for ship detection.

Adaptive threshold algorithms are designed to search for pixel values which are unusually bright compared to those in the surrounding area. This is done by setting a threshold which depends on the statistics of the surrounding area (hence the term adaptive). Pixel values which lie above the threshold are declared unusually bright and therefore likely to be samples from a target. The statistics of the surrounding area are usually analysed by taking a ring of image samples around the pixel under test. The most basic setup is shown in Figure 1. Here, the pixel under test is surrounded by a guard "ring" and then a background "ring". The purpose of the guard "ring" is to ensure no pixels of an extended target are included in the background "ring" and hence that the background "ring" is representative of the underlying background statistics. All windows are moved one pixel at a time across the whole image. Obviously, many variations of this basic setup

Figure 1: Typical window setup for an adaptive threshold detector

are possible. The exact relationship between the threshold and the background statistics depends on which adaptive threshold algorithm is being considered.

A natural aim in designing adaptive threshold detectors is to ensure that the probability of a false alarm is constant. In this case, the threshold is chosen so that the percentage of background pixel values which lie above the threshold is constant. If this is done then the false alarm rate (number of false alarms per unit area of imagery) is likewise constant. Consequently, such detectors are called constant false alarm rate (CFAR) detectors. One method for CFAR detection is to work directly with the histogram of the background values and set the threshold at the appropriate point in its tail. However, the required false alarm rates are usually very low and large background samples with an attendant computational burden would be required to accurately estimate the threshold. Thus, it is more usual to model the background distribution parametrically and to use the background samples to estimate the model parameters.

Choosing a parametric distribution model for the background is equivalent to specifying the associated parametric probability density function $f(x)$, where x ranges through the possible pixel values. Once $f(x)$ has been chosen and its parameters estimated from the background samples, the probability of false alarm (PFA) for the threshold T is given by

$$\text{PFA} = 1 - \int_{-\infty}^{T} f(x)dx = \int_{T}^{\infty} f(x)dx. \tag{10}$$

Thus, designing a CFAR detector involves solving the equation (10) for the threshold T in terms of the specified PFA and the estimated parameters of the pdf $f(x)$. An analytic solution to this problem is not always possible and numerical methods may be needed. One approach is to search for the correct value of T by trial and error. This is fairly straightforward since cumulative distribution functions are increasing. An example of how this is done is discussed at the end of the next section when we cover the work of Jiang *et al.* in [54]. Another approach which may be more efficient is as follows. Since the

CFAR detector takes the form

$$x_t > T \qquad \Longleftrightarrow \qquad \text{target}$$

where x_t is the pixel value under test and since the cumulative distribution function is increasing, an equivalent detection test is

$$\int_{x_t}^{\infty} f(x)dx < \text{PFA} \qquad \Longleftrightarrow \qquad \text{target.} \qquad (11)$$

In either approach, the integral $\int_{x_t}^{\infty} f(x)dx$ or equivalently $\int_{-\infty}^{x_t} f(x)dx$ only needs to be evaluated numerically.

A commonly used statistical model is the Gaussian distribution. This is often applicable since the central limit theorem states that the average of a large number of identically distributed random variables tends to have a Gaussian distribution. Designing a CFAR detector based on the Gaussian model leads to the ubiquitous two parameter CFAR detector. In this case, the detector is

$$x_t > \mu_b + \sigma_b t \qquad \Longleftrightarrow \qquad \text{target} \qquad (12)$$

where x_t is the pixel value under test, μ_b is the background mean, σ_b is the background standard deviation, and t is a detector design parameter which controls the PFA (or equivalently the false alarm rate). In practice, t is fixed and μ_b and σ_b are estimated from the samples in the background "ring". The exact relationship between t and the PFA is given by

$$\text{PFA} = \frac{1}{2} - \frac{1}{2}\text{erf}\left(\frac{t}{\sqrt{2}}\right). \qquad (13)$$

Given the desired value of PFA, this must be solved numerically for the corresponding design parameter t.

Although the Gaussian pdf is commonly used, it is not an accurate model of radar imagery unless a large number of looks have been averaged. More appropriate models (although still not perfect) for radar intensity imagery are the negative exponential for single look imagery and the gamma distribution (with fixed order parameter) for multi-look imagery. Designing CFAR detectors for these distributions leads to the well known Cell Averaging CFAR detector. In this case, the detector is

$$x_t > \mu_b t \qquad \Longleftrightarrow \qquad \text{target} \qquad (14)$$

where x_t is the pixel value under test, μ_b is the background mean and t is a detector design parameter which controls the false alarm rate. As before, t is fixed and μ_b is estimated from the samples in the background "ring".

It is of interest to consider here imagery with a constant normalised variance where by **normalised variance** we mean the variance divided by the mean squared. For such imagery, the standard deviation is proportional to the mean and so the two parameter CFAR detector (12) effectively becomes a Cell Averaging CFAR detector but with a different design parameter t. Both the negative exponential and the gamma (with fixed order parameter) distributions have constant normalised variance. In fact, the normalised variance is equal to the reciprocal of the order parameter. It follows that the two parameter

CFAR detector is CFAR when applied to exponential and gamma distributed data. In [135], Wackerman *et al.* provide an interesting figure which compares how the FAR varies with t for the two parameter CFAR detector when applied to Gaussian, exponential and gamma distributed data. As we discuss in Section 5.4, Wackerman *et al.* also consider the K-distribution in this figure.

A useful variation of the adaptive threshold detector involves searching for groups of pixel values which are unusually bright. In this case, the central single pixel target window is expanded to incorporate more than one pixel. The mean value over the target window, μ_t, is then calculated and compared with the threshold. Thus the two parameter CFAR detector becomes

$$\mu_t > \mu_b + \sigma_b t \qquad \Longleftrightarrow \qquad \text{target} \qquad (15)$$

However, care must be taken here since the statistics of the background are calculated from individual pixels and hence are no longer directly comparable to those of μ_t. Fortunately, a simple adjustment corrects this problem. Averaging n samples from a Gaussian distribution with mean μ and standard deviation σ produces another Gaussain distribution with mean μ and standard deviation $\sigma/\sqrt{n}$. Therefore, if the target window consists of n pixels, the detector (15) is still CFAR but the design parameter t needs to be replaced with $t/\sqrt{n}$ to achieve the same false alarm rate.

A similar argument works for the cell-averaging CFAR detector. Averaging a negative exponential or gamma distribution over n samples produces another gamma distribution with the same mean (but different order parameter) and so the cell-averaging CFAR detector remains a CFAR detector when multiple pixel target windows are used. In this case the detector becomes

$$\mu_t > \mu_b t \qquad \Longleftrightarrow \qquad \text{target} \qquad (16)$$

However, since the order parameter is changed by the averaging process, the adjustment of the design parameter t to achieve the same false alarm rate is no longer simple. Oliver and Quegan, [86], note that

$$\text{PFA} = \frac{\Gamma(L,t)}{\Gamma(L)}. \qquad (17)$$

where L is the order parameter of the gamma distribution. Again, this must be solved numerically to find the correct value of t for a given PFA. The value of L is the effective number of looks produced by the averaging process.

We emphasise that the detectors, (12), (14), (15) and (16) are more generally applicable than we have so far implied. There are several background distributions for which they are CFAR. For instance, Goldstein, [42], shows that (12) is CFAR for log-normal and Weibull clutter. But more to the point, even if the detectors are not CFAR for the true background distribution of the data, they will still only pick out bright pixel values. This has been noted in particular by Novak *et al.* in the paper [84] where the detector (12) was observed to be an effective prescreener for SAR imagery despite not being CFAR for the imagery concerned. Since CFAR algorithms are often applied in situations where they are not truly CFAR we believe it is probably more accurate to refer to them as adaptive threshold algorithms.

An important consideration in using these detectors is choosing the design parameter t. As noted, the aim is usually to achieve a desired false alarm rate. However, the

relationship between t and the false alarm rate can be complex and setting t may require difficult calculations. Often the approach taken, and especially when the detectors are not CFAR for the imagery concerned, is to set t empirically. Thus a range of values of t is tried on real imagery and the best value selected on the basis of performance. This has the added benefit of allowing the probability of detection to be taken into account too. We will say more about setting t when specific ship detection algorithms are discussed.

There are many possible variations of the size and shape of the target, guard and background windows and how they are moved across the image. The final choice will be determined by the detection problem at hand but will also need to take account of computational burdens. The standard practice in ship detection is to use square windows and move them as a unit in single pixel steps. Square windows are the common choice because they allow computational efficiency and there is a lack of *a priori* knowledge of preferred shapes and orientation for ships and ocean clutter. Wackerman *et al.* expand upon this point in [135]. The sizes of the windows should be related to the size of the ships to be detected and the resolution of the radar. The target window should be a similar size to the smallest ship to be detected, the guard window should be a similar size to the largest ship of interest and the background window should be large enough to allow good estimates of the background statistics without being so large as to include other targets or non-background objects. Again, we refer the reader to [135] for a discussion of these considerations.

5.3   Adaptive threshold algorithms for ship detection

The work by Wackerman *et al.*, [135], and the related work by Freidman *et al.*, [38] and [39], is some of the most convincing work on ship detection in the literature. This algorithm was developed by ERIM International Incorporated, [94]. (ERIM is an acronym for the Environmental Research Institute of Michigan) We will begin with a discussion of their algorithm.

Wackerman *et al.* [135] use the two parameter CFAR detector with a multiple pixel target window, (15), for detecting ships in Radarsat-1 C-band SAR imagery. They mainly consider the ScanSAR Wide B mode with 200 m resolution and 100 m pixel spacing although finer resolution imagery is also used. They note that their detector is CFAR for both Gaussian and exponentially distributed clutter but that more appropriate models for ocean clutter are the gamma distribution and the K-distribution. They point out that the false alarm rate for these distributions depends only on the normalised variance (variance divided by the mean squared) and so the detector will be CFAR for them as long as the normalised variance is constant across the image.

Wackerman *et al.* use standard target, guard and background windows and they include a comprehensive discussion of how and why they choose the final shapes and sizes. The target window is 3x3 pixels, the guard window is 5x5 pixels and the background window is 13x13 pixels. One of their considerations was severe processing time constraints which may explain why the guard and background windows are smaller than the 7x7 pixels and 19x19 pixels windows quoted in [94]. Also, for the finer resolution imagery (100 m resolution, 50 m pixel spacing) all the window sizes are approximately doubled: the target window is 5x5, the guard window is 7x7 and the background window is 25x25 pixels. Note

also that Wackerman *et al.* suggest that if more time were available, better results may be obtained by using a dynamic target window. By this they mean allowing the target window to change size and shape and thereby incorporate all the target energy. One way this could be done is to use a series of target windows of varying size and shape and choose the one with largest mean μ_t.

Wackerman *et al.* stress the importance of setting the design parameter t correctly and it is clear from their work that considerable effort and analysis was spent in doing so. A very informative figure is presented showing how the probability of false alarm varies with the design parameter t for the various background distributions considered. This figure emphasises the importance of using the correct background distribution model when setting t theoretically. In practice however, Wackerman *et al.* set t empirically using imagery which was ground-truthed by eye. The theoretical and empirical results are compared and several interesting observations are made. Wackerman *et al.* found that the best trade-off between probability of detection and false alarm rate occurred when $t = 5.5$. Note that in a related report, [39], $t = 5.0 - 5.5$ is recommended for the low resolution imagery and $t = 5.5 - 6.0$ is recommended for the high resolution imagery.

In an interesting aside, Wackerman *et al.* compare their empirical results with the theoretical results and conclude that their observed false alarms rates are consistent with either a gamma distribution with normalised variance between 1 and 4 or a K-distribution with a normalised variance of 2. A significant point arises in making this comparison. It relates to the adjustment of t in (15) to take account of the effect of μ_t being an average over the target window. Wackerman *et al.* point out that for Radarsat imagery the pixel spacing is different to the resolution and hence the number of pixels in the target window is different to the number of independent samples. It is the latter quantity which is important when adjusting t.

Wackerman *et al.* designed their detector for application to Radarsat-1 ScanSAR Wide B Mode imagery (200 m resolution, 100 m pixel spacing). Detection performance was analysed using SAR imagery of a fishery. Only partial ground truth information was available and visual analysis was used to oversee performance assessment. It was known that 272 ships were present within the fishery. The ship lengths were also known. The imagery also included areas of open ocean which were known to contain no ships. Through clever use of this partial information, Wackerman *et al.* were able to assess detection performance. The probability of false alarm was estimated using the open ocean areas and it was assumed to be the same in the fishery. This allowed the number of false alarms in the fishery to be predicted. Consequently, the number of true detections (hits) could be estimated by subtracting the predicted number of false alarms from the total number of detections. It was assumed that the missed detections were the smaller ships. The resulting estimate of performance was that all ships over 35 m long could be detected with a false alarm rate of 0.01% (i.e. the probability of false alarm is 0.01). The same scene was also available at higher resolution (100 m resolution, 50 m pixel spacing) in which case the estimate of performance was that all ships over 32 m long could be detected with a false alarm rate of 0.002%. Note that doubling the resolution does not halve the minimum detectable ship length.

Eldhuset, [32] and [31], also uses the detector (15) for ship detection in SAR imagery. His application is to both SEASAT and ERS-1 SAR imagery. Time constraints were

a consideration and for this reason Eldhuset has an unusual arrangement for moving the detector's windows across the imagery. The target window is 2x2 pixels, the guard window is 10x10 pixels and the background window is 20x20 pixels. The guard and background windows are moved as a unit in steps of 10 pixels across the image. At each location of these windows, the background statistics are estimated and the detection threshold set. The target window is then moved around within the guard window and at each pixel the detection test (15) is applied.

Actually, rather than (15) Eldhuset uses the equivalent test $4\mu_t - 4\mu_b > q\sigma_b$ where q is a fixed design parameter. Further, Eldhuset does not actually estimate σ_b from the background pixel values, but rather uses the well-known approximation for intensity SAR imagery that $\sigma_b = \mu_b/\sqrt{N}$, where N is the effective number of looks, and an analogous approximation for amplitude imagery. We note here that using such estimates for σ_b means that Eldhuset is in effect using the detector (16).

Eldhuset recommends setting $q = 5.0$. He provides no justification for this choice. Presumably it was set by trial and error. Setting $q = 5.0$ translates to setting $t = 1.25$ in (15) which is equivalent to $t = 2.5$ in (12). This is lower than the value of 5.5 recommended by Wackerman *et al.* [135]. The discrepancy can probably be explained by the fact that Eldhuset's system includes further stringent tests to reduce the false alarm rate. We will say more about Eldhuset's tests in later sections.

Eldhuset states that the background and guard windows are moved in multi-pixel steps for reasons of computational efficiency. He adds that this means the detection results will vary depending on the exact starting point of the windows but notes that this is only a problem in areas of the imagery that are inhomogeneous and the results are robust even so. We note here that the computational methods described by Wackerman *et al.* [135], are very efficient and we believe there is little to gain by Eldhuset's approach.

Another group to use the adaptive threshold algorithm (15) is Casasent *et al.* [19]. Their application is to SIR-C SAR imagery with a resolution of 50 metres. Ground truth was obtained by eye from higher resolution (12.5 metres) imagery. They use a standard setup with square windows. By measuring performance on their substantial database, they determine acceptable window sizes and design parameter value. They recommend a 2x2 target window, an 8x8 guard window, a 31x31 background window and they set $t = 5.0$ in (15). However, they note a high false alarm rate in calm seas and so make an *ad hoc* modification to the basic detector. They force a lower bound on the value of σ_b — if σ_b is less than the lower bound it is reset to the lower bound. Unusually, they define the probability of false alarm as the number of false alarms divided by the number of targets and so care is needed in interpreting their results. They also mistakenly claim that the idea of using a guard window is new (which is possibly explained by a background in IR imagery). Their results are also puzzling in that the best setting for t does not seem to scale with the size of the target window as theory would dictate — see in particluar section 3.2.2. However, they do verify that a guard window improves detection performance and they examine the question of whether the target pixel variance σ_t aids detection performance by considering the detector $|\mu_t - \mu_b| + \sigma_t > \sigma_b t$ but without success.

In a final section of the paper [19], Casasent *et al.* propose some modifications to improve the performance of the detector. They pay special attention to three difficult detection cases: ships close to land; several ships in close proximity; ships with large side

lobes. Casasent *et al.* comment that in all these cases, μ_b and σ_b increase, and so the detection threshold in (15) is too high resulting in missed detections. For the first case, they simply use a land mask and remove land pixels from the background "ring". For the second case, they suggest first removing pixel values x from the background "ring" which satisfy $x > \mu_b + 2.0\sigma_b$ and then recalculating μ_b and σ_b. This obviously also helps avoid missed detections in the third case. There is an additional consideration in the third case — false alarms in the sidelobes. To remove these, a filter is used to detect and then mask out the sidelobes of detected targets. This last modification has the additional benefit of yielding a more accurate final shape for the detected ship itself. Casasent *et al.* then go on to propose a set of rules aimed at reducing the false alarm rate when the background is very noisy, very dim or very smooth. The rules are *ad hoc* and mostly involve adjusting the value of $\sigma_b t$ in (15) based on the size of μ_b and σ_b.

Next we describe the adaptive threshold detection algorithm used in the Alenia Areospazio ship detection system developed by Ferrara *et al.*, [33], [34] and [35]. The application is to SIR-C/X-SAR L band SAR imagery, initially single look complex (SLC) with a resolution of 4 metres but they also process the imagery to a range of multi-looks with correspondingly different resolutions. Unfortunately, the original description of the system [35] lacks full details. There, they merely state that the detector is a threshold algorithm and that the threshold is calculated from a user specified value of the FAR. The more recent descriptions, [33] and [34], are more comprehensive with full details of the algorithm given in [146]. They use a hybrid algorithm involving both a cell-averaging CFAR detector (CA-CFAR) and an order statistics CFAR detector (OS-CFAR). The CA-CFAR detector has been discussed above and uses the detection rule (14). The parameters recommended in [146] for 4 look data are a 1x1 target window, a 3x3 guard window, a 7x7 background window and $t = 5.65$. The order statistics CFAR detector is identical except that μ_b in (14) is replaced with an order statistic and $t = 4.63$. The particular order statistic which replaces μ_b is not specified in [146] but presumably it is the median value. It is noted in [146] that the CA-CFAR performs well in homogeneous environments while the OS-CFAR performs well in inhomogeneous environments (such as when interfering targets are present in the background window). Hence, the strengths of these two detectors are combined by calculating a measure of the homogeneity in the background window and then applying the appropriate detector. The homogeneity measure used is the ratio σ_b/μ_b. When $\sigma_b/\mu_b < 1/\sqrt{L}$ where L is the number of looks, the local area is assumed to be homogeneous and the CA-CFAR is applied, otherwise the OS-CFAR is applied.

In [33], Ferrara *et al.* provide a succinct schematic diagram of the hybrid detector just described. However, they imply in the introduction that the design parameter t in (14) is not fixed at the values specified in [146] but rather is calculated from user specified false alarm rates using the appropriate background pdf. Details of the pdf are not given but a good choice is obviously the gamma distribution as mentioned in [146]. Ferrara *et al.* note that the detector has improved performance over their previous detector and especially in the case of small targets in low resolution imagery. A more detailed analysis is given by Ferrara and Torre in [34]. There, they model the background with a chi-squared pdf and use that for selection of the design parameter t. They also model the target with a Rice distribution and hence can assess probabilities of detection. They consider a range of different resolutions corresponding to multi-look data of 1, 4, 9, 16 and 64 looks. They

report a theoretical probability of detection of greater than 80% except for the 64 look data where PD drops to less than 50%. Interestingly, they consider the case where isolated detected pixels are considered too small to be ship pixels and are discarded (of course the validity of this depends on image resolution and ship size). In this case they use a different FAR (presumably higher) and hence get better probabilities of detection.

Cusano *et al.*, [25], describe a ship traffic monitoring system which operates at ESA's European Space Research Institute on ERS 1/2 Quick Look SAR imagery. At its heart is an adaptive threshold algorithm for detecting ship-like pixels. They comment that this is an improved detector which has been designed to cope with the high variability of sea surface scattering. The core of the detector is a multi-target pixel cell averaging CFAR detector (16). However, it has been improved by allowing the design parameter t to be *a function of the variability of the SAR image level*. Cusano *et al.* do not elaborate on exactly how t is adjusted neither do they specify detector window sizes and shapes. Fortunately, Lombardo and Sciotti, [70], provide a brief description of the detector which allows some insight to the adjustment of t. They state that the technique *scales directly the detection threshold as a function of variance*. Cusano *et al.* add that areas of very high contrast are automatically detected and masked out. Their results show that detector works well and is not confounded by closely spaced ships or the presence of ship wakes.

The paper, [73], on ship detection in Radarsat imagery by Lopes *et al.* is brief and poorly written. Consequently it is hard to assess. They mention using a likelihood ratio test, however, the core of their detection algorithm appears to be the cell-averaging CFAR detector (16). Their setup has no guard window and only consists of a multi-pixel target window which is contained in a larger background window. Their detection criterion is $\mu_t/\mu_b > t$ for some fixed threshold t. Interestingly they derive the pdf of the ratio μ_t/μ_b under the assumption that the target window only contains background pixels and the background distribution is a negative exponential (which is consistent with a single-look intensity image). From the pdf, they obtain an expression for the PFA in terms of t. They note that when clutter samples are correlated, the arithmetic mean of the intensity is not a minimum-variance estimator of power levels and hence they suggest the use of a spatial whitening filter (SWF). The SWF was introduced in [63] and shown to be a superior method of speckle reduction compared to noncoherent spatial averaging.

The system proposed by Rye *et al.* [107] for ERS-1 data, uses a fixed global threshold t to detect ship pixels. Importantly, this is combined with a preprocessing step which normalises the imagery by dividing each pixel value by its local mean μ_b (calculated over a window of 30x30 pixels). Thus the detection test is $x_t/\mu_b > t$ which is just a rearrangement of (14) and in effect Rye *et al.* are using a cell-averaging CFAR detector. However, since this detector uses no guard window, its performance would be expected to be worse than a traditional cell-averaging CFAR detector. Rye *et al.* suggest several ways to improve performance: detection only if N pixels are above the threshold in an area of M pixels; searching for adjacent pixels above the threshold; and using two thresholds with pixels above the lower threshold being used to link together pixels above the upper threshold.

The ocean monitoring workstation (OMW) is a high profile ship detection system which we discuss in the next section. Here, we mention the detection algorithm used by its precursor - the ocean features workstation (OFW). Rey *et al.* report in [103] that the OFW used a two parameter CFAR detector with design parameter $t = 5.0$. They provide

few other details. However, they do note that while the detector was adequate for ERS-1 SAR data, it did not perform so well on Radarsat-1 data - the false alarm rate was too high. The reason proffered being that Radarsat data is higher resolution. They found that a CFAR detector based on a K-distribution model for the ocean clutter gave better performance.

The CFAR detector used by Jiang *et al.* in [54] is of interest in that it is a CFAR detector which is based on a non-parametric estimation of the ocean clutter pdf. Thus they use equation (10) where $f(x)$ has been estimated non-parametrically. This work grew out of earlier work with a K-distribution based CFAR detector, [55], where it was found that the K-distribution was not always a good fit to the data. The non-parametric technique they use is called probabilistic neural networks but it is essentially a Parzen window method of approximating the pdf from histogram data. The approximation is

$$f(x) = \sum_{i=1}^{M} P[i]G(i,\sigma)(x)$$

where M is the number of image grey levels, $P[i]$ is the normalised histogram and

$$G(i,\sigma)(x) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(x-i)^2}{2\sigma^2}\right).$$

Computing pdfs this way can be computationally expensive. Jiang *et al.* solve this problem by using look-up tables for calculating the quantities $G(i,\sigma)(x)$. The other important issue in Parzen window methods is choosing the correct kernel "width", σ. Jiang *et al.* do this empirically with the use of training and validation data. Their method also involves the definition of an error function for comparing the difference between two histograms.

Jiang *et al.* do not mention their CFAR window setup. However, it is implicit that, as in the work by Vachon *et al.* [127] to which they refer, each image is divided into frames and each frame is processed as a whole. Thus the pdf is estimated for the whole frame and the resulting threshold T is in turn applied to the whole frame. While this means the algorithm is more like a global threshold algorithm rather than an adaptive one, it should effectively be adaptive as long as the frame size chosen is small enough so that the background statistics are stationary across the whole frame.

Having specified how to estimate the pdf $f(x)$, Jiang *et al.* then discuss how to solve (10) for the threshold T given the PFA. This is an iterative process which relies on the fact that the distribution function $F(x) = \int_{-\infty}^{x} f(t)dt$ is increasing. The first step is to find the integer T_0 with $F(T_0) < 1 - PFA < F(T_0 + 1)$ in which case $T_0 < T < T_0 + 1$. Next the value $T_1 = T_0 + 1/2$ is tried and $F(T_1)$ is compared with $1 - PFA$ to determine if $T_0 < T < T_1$ or $T_1 < T < T_0 + 1$. This process then continues with the interval containing T being halved each time until the desired accuracy is reached. (We note here that according to standard optimisation theory it is more efficient to conduct this search by partitioning the interval containing T into golden ratio proportions or better yet using a Fibonacci search rather than halving it every time.) Calculating $F(x)$ for given values of x is still non-trivial and in the case being dealt with here, a rational approximation to the error function erf is used.

Jiang *et al.* compare the performance of their algorithm with a K-distribution based CFAR detector. The actual detector is not specified but given their earlier work, [55], and

the reference to that of Vachon *et al.*, [127], it is likely that the comparison was with the OMW detector discussed in the next section. The results show that the method of Jiang *et al.* is faster and models the data better.

5.4   The OMW detector

The ocean monitoring workstation (OMW) is one of the few commercially available software packages for the detection of ships in SAR imagery. The OMW has other important capabilities as well, such as oil spill monitoring and wind and wave field extraction, but we will only be interested in ship detection here. Its main application is to near real-time analysis of Radarsat and ERS imagery (but it accepts any CEOS Level 1 SAR imagery). It has a long history of development and represents state-of-the-art technology. It was developed by Satlantic Inc. under contract to Fisheries and Oceans Canada. Funding for the contract comes from several Canadian government agencies including the Department of National Defence (DND) and the Canada Centre for Remote Sensing (CCRS). DND's interest, through Defence Research Establishment Ottawa (DREO), is in ship detection for wide area surveillance. DREO has provided substantial technical input into the OMW.

Information on the OMW and it's ship detection algorithms may be found on Satlantic's web page [7], on Canadian Centre for Remote Sensing's (CCRS) web page [8], and in the papers, [126], [47], [49], [134], [133], [102], [14], [48], [128], [127] and [103]. Not much has appeared in the literature recently concerning ship detection with the OMW so it is unclear what the current state of play is. In their summary in 2000 Vachon *et al.* [126] note that the research efforts at DREO then were directed towards reducing the FAR of the OMW and modifying the algorithms for application to fully polarimetric data (in anticipation of the launch of Radarsat-2).

Being a commercial product, it would not be surprising if full details of the OMW procedures have not been disclosed and obviously we can make no guarantee that the description in this report is complete and up to date. However, as mentioned, the development of the OMW was a collaborative effort and many details have appeared in the literature. The most detailed and comprehensive reports are in [49] and [102]. In particular, in section 2.0 of the report [102], Rey *et al.* provide a summary of the OMW ship detection algorithm at the time and go on to suggest modifications in section 5.0. This is the most complete description available in the open literature.

The OMW ship detection algorithm is essentially a CFAR algorithm based on modelling the ocean clutter with an L-look K-distribution. It is well known that the K-distribution is a good model for single point clutter statistics in homogeneous regions of high resolution radar imagery, see for instance [9]. Detailed research, as discussed below, has confirmed that it is appropriate for satellite borne SAR such as Radarsat and ERS. However, we also warn that it is not appropriate for all SAR imagery. For instance, it was reported by Yeremy *et al.* in [143] that the OMW ship detection algorithm could not be applied to SAR data collected by the CCRS C/X SAR because the nadir region of the imagery did not conform to the K-distribution model.

---

[7]http://www.satlantic.com

[8]http://www.ccrs.nrcan.gc.ca

The pdf for the K-distribution of intensity imagery is

$$f(x) = \frac{2}{x} \left(\frac{L\nu x}{\mu} \right)^{\frac{L+\nu}{2}} \frac{1}{\Gamma(L)\Gamma(\nu)} K_{\nu-L} \left[2 \left(\frac{L\nu x}{\mu} \right)^{\frac{1}{2}} \right] \qquad (18)$$

where L is the number of statistically independent looks, ν is an order parameter, μ is the mean value, $\Gamma(\cdot)$ is the gamma function, and $K_{\nu-L}(\cdot)$ is the modified Bessel function of order $\nu - L$. A similar pdf can be derived for amplitude imagery but intensity imagery is usually used. Vachon *et al.* [134] note (without elaboration) that the order parameter ν models the wind variability within the ocean scene.

The number of statistically independent looks L is determined by the SAR processor and is presumed to be known. Hence constructing a CFAR detector based on the K-distribution involves fitting the distribution (18) to the ocean clutter by estimating the mean μ and the order ν from samples and then determining the detection threshold T by solving (10) for the estimated parameters and the given PFA. Both these tasks are difficult and have been the subject of extensive research. We discuss the different methods later in this section.

The OMW does not use the standard CFAR detector window setup in Figure 1, rather it divides the image into frames (small imagettes), [134], [49], [102] and [103]. The ocean is assumed to be homogeneous within each frame and the whole frame is used to estimate the parameters μ and ν of the K-distribution (18). Once the threshold T has been determined it is applied to the whole frame. It is presumed that this approach has been chosen for computational efficiency — determining the threshold T is numerically intensive and it is not practical to do so for each pixel individually. Selecting the appropriate frame size is thus a trade off between accuracy and efficiency.

Not many details are readily available on frame size specification for the OMW. It is noted in [134], that recommendations have been made in the CCRS/AD/ADS internal report, [16], but that report was not obtained. Henschel *et al.* [49] state that a configuration file is used to specify the frame size and that it is Radarsat beam specific. However, the only other comment they make on the topic is that frame size is significantly smaller than 3200 pixels square. The description by Rey *et al.* in [102] is likewise scant on detail. They merely note that the frame size is typically 256 pixels square. In older work, Rey *et al.* [103] state that a rule of thumb derived from experience with ERS-1 data is that stationarity of ocean statistics can be assumed over no more than a 10km x 10km region. They recommend using frames no larger than 1000 x 1000 pixels for Radarsat data. They go on to say that non-stationarity may occur in smaller regions and care is needed.

We note that the absence of a guard window means that target pixels can corrupt the process of fitting the K-distribution. Rey *et al.*, [102], state that the OMW partially alleviates this problem by flagging any saturated pixels as possible targets and hence excludes them from the background. They go on to suggest a possible solution in the case of non-saturated targets is to do a first pass with an empirically derived threshold T which detects high probability targets and then remove them from the background samples.

We return now to the question of how to solve (10) for T assuming that μ and ν have already been estimated. Due to the intractability of the K-distribution, this needs to be done numerically. Note also that, since the same threshold T is applied to the whole frame,

it is more efficient to solve (10) for T rather than use (11). The best reference as to how the former is actually done in the OMW is the report, [102], by Rey *et al.*. They state that *T was originally calculated from the order parameter for a given number of statistically independent looks, L, and CFAR value using a look up table which defined the parameters of an approximation derived from numerical calculations.* They go on to state that *in order to refine the accuracy of the approximation, the look up table method is now used for an initial value, and the iterative method based on the analytic expression developed at DREO, [103], is used to calculate the final value of the K-distribution critical intensity threshold.* It seems likely that the original look-up table method is the one described by Campbell *et al.* in [14]. Presumably, the iterative technique referred to here is the same as the search technique used by Jiang [54] which we discussed at the end of Section 5.3.

The standard method of parameter estimation is the maximum likelihood solution. However, Blacknell, [9], notes that this is analytically intractable for the K-distribution and suggests three moment based alternatives. Based on Blacknell's results, two of these alternatives have been closely examined for the OMW. The first uses the mean and variance of the data and is referred to as the MV method. It provides the estimates $\tilde{\mu}$ and $\tilde{\nu}$ of the mean and order parameter μ and ν by solving the equations:

$$\tilde{\mu} = \langle x \rangle \tag{19}$$

$$\left(1 + \frac{1}{\tilde{\nu}}\right)\left(1 + \frac{1}{L}\right) = \frac{\langle x^2 \rangle}{\langle x \rangle^2} \tag{20}$$

where $\langle \cdot \rangle$ is the expectation operator. The second estimator uses the mean of the data and mean of the log of the data, and is referred to as the MML method. It requires the equations:

$$\tilde{\mu} = \langle x \rangle \tag{21}$$

$$\ln(\tilde{\mu}) - \psi(\tilde{\nu}) = \ln(\langle x \rangle) - \langle \ln(x) \rangle + \psi(L) - \ln(L) \tag{22}$$

to be solved where ψ is the digamma function.

Early work by the CCRS and DREO researchers, [14] and [127], verified the suitability of the K-distribution together with the MV method of parameter estimation for modelling ocean clutter in both ERS and Radarsat SAR imagery. We mention in particular here the Pearson diagrams in [127] where ERS and Radarsat ocean clutter are seen to be K-distributed. It seems likely that at the time, the OMW algorithm was based on the work of Campbell *et al.* as documented in [14].

Concurrent with the early work by the CCRS and DREO researchers, Rey *et al.* [103] took a more in depth look. They considered both the MV and the MML methods of estimation and used both the χ^2 test and the Kolmogorov-Smirnov tests to test the goodness of fit between the sample probability distribution and the estimated K-distribution. Blacknell's third estimation method, [9], was not considered since his analysis showed that it was generally less accurate than the MV and MML methods. The experiments were conducted on ERS-1 data. While the results were not entirely conclusive they were consistent with the theoretical analysis of Blacknell, [9], and Oliver, [87]. On this basis, it was concluded that the MML approach is better for low values of ν and the MV approach is better for high values of ν. The empirically obtained cross-over point being

$$\nu = 6.1L + 1.25 \tag{23}$$

where L is the number of effective looks.

Based on the results in [103], Rey *et al.* conclude that the K-distribution is a good model for ocean clutter. They recommend that both the MV and the MML methods be used for order parameter estimation and that (23) be used to determine which estimation method is actually applied. They go on to outline the steps of a complete CFAR detection algorithm. They report that this algorithm has better performance than the old two-parameter CFAR one. The new algorithm includes a clever idea for computational efficiency which we mention here. Rather than proceeding straight to solving (10) for T given the user specified probability of false alarm (which is computationally intensive), the idea is to use (10) to calculate the probability of false alarm when T is set to the maximum pixel value in the frame. If this probability is less than the user specified one then the frame has no pixels which will be declared targets by the detector and processing is stopped.

It is noted by Rey *et al.*, [102], and Henschel *et al.*, [49], that sometimes when fitting the K-distribution to Radarsat data, the estimation order parameter ν is negative. The implication is that the K-distribution is not a good model for such data. Consequently the OMW was modified to use a χ^2-distribution look-up table in these cases. Rey *et al.* [102] write the χ^2-distribution with $2n$ degrees of freedom as

$$f(x) = \frac{x^{n-1}}{2^n \sigma^{2n} \Gamma(n)} \exp\left(-\frac{x}{2\sigma^2}\right) \qquad \text{for } x \geq 0 \text{ and } 0 \text{ elsewhere} \tag{24}$$

where n is the order parameter and σ is the standard deviation of the distribution. Rey *et al.* provide a discussion of how the χ^2-distribution relates to the K-distribution and go on to justify the use of the χ^2-distribution when the estimated value of ν is negative. They note that the χ^2-distribution, under a suitable change of variables, becomes the gamma distribution (2) which is, in turn, the limit as ν approaches infinity of the intensity K-distribution. Thus, in theory, the order parameter n is equal to the effective number of looks L. Rey *et al.* [102] state that at the time of their report, the OMW used a χ^2-distribution with $n = L$ whenever the MV estimate of the K-distribution order parameter ν was negative. However, they note that there is uncertainty in the value of L and that n can also be estimated by

$$n = \frac{\langle x \rangle^2}{\langle x^2 \rangle} \tag{25}$$

where $\langle \cdot \rangle$ is the expectation operator. In [102] they report good results using this estimate for n and consequently recommend the OMW be modified accordingly. We note here that setting the detection threshold T when using a χ^2-distribution involves iteratively solving (10) for the given PFA in much the same way as when using the K-distribution.

The initial aim of the report [102], Rey *et al.*, was a further investigation of which of the two K-distribution parameter estimation schemes gives better results. However, the scope was extended to include an investigation of the χ^2-distribution as an alternative ocean clutter model when it was found that K-distribution did not always fit the data. The appropriate choice of the probability of false alarm (CFAR value) was also considered with a range of CFAR values being tested in order to take account of probability of detection as well as probability of false alarm. Radarsat data from the Scotia Shelf exercise in 1996 and the MARCOT 1997 naval trial was used. The ground truth was provided by

expert human imagery analysts. It was found that K-distribution failed to fit most of the MARCOT scenes. While this could not be fully explained some suggestions were made as to why. The χ^2-distribution was shown to be a suitable alternative. It was found that for a given CFAR value, the MML algorithm generates more false alarms than the MV algorithm but less than the χ^2 algorithm. But it was also found that the MML algorithm detects more targets too and hence gives a better trade off between detecting targets and rejecting false alarms. The final recommendation is to use the MML based K-distribution with a CFAR of 10^{-8}. It was also noted that the χ^2-distribution with a lower CFAR value performed well too but more work is needed to confirm this.

The MARCOT naval exercise off the East Coast of Canada in 1998 provided an extensive data set for development and validation of ship detection algorithms. The idea of this trial was to conduct an operational evaluation of the OMW. However, Henschel *et al.*, [49], also used the trial data to investigate the best procedure for driving the OMW ship detection algorithm. Their work appears to be direct extension of that of Rey *et al.*, [102]. They address similar questions and they reach similar conclusions. The investigation by Henschel *et al.* was based on preliminary data since the full ground-truthed statistics were not available at the time. Unfortunately, an analysis of the complete MARCOT'98 data has not yet appeared in the open literature.

This completes our review of the OMW ship detection algorithm except for the following miscellaneous notes. The OMW manual, [111], states that when used with Radarsat Fine beam mode the OMW can detect ship of 5–10 m length in low wind. Campbell *et al.*, [14], note that taking into account the multi-pixel extent of targets can improve their detectability although they provide no details on how to do this for OMW-type algorithms. Also, in the OMW algorithm description of Rey *et al.*, [102], it is mentioned that: Radarsat imagery has dark strips which are removed in a preprocessing step; if there is land within a frame then its statistics are replaced with those of the nearest neighbouring frame which does not contain land; detected pixels are clustered into groups of contiguous pixels; and any candidate targets which are too large, too small, too close to land or too close to another target are eliminated.

We conclude this section with some comments on other work with K-distribution based CFAR detectors. The first we mention is by Jiang *et al.*, [55]. This work is closely related to the OMW. Jiang *et al.* recommend testing the goodness of fit of the K-distribution (by a mean squared error criterion) and only using it when the error of the fit is below a pre-set value. They also use two detection thresholds — the first generates a very low probability of false alarm while the second is less stringent. The idea is that the second threshold picks up less bright pixels missed by the first threshold. Since these pixels are less likely to be ship pixels, Jiang *et al.* subject them to a wake test. Only those pixels for which the wake test is positive are retained.

Robertson *et al.*, [105], have also used a K-distribution based CFAR detector. However, unlike the OMW, their detector uses a standard roving window to determine local clutter statistics which are then fitted to a K-distribution. No comment is made on the computational cost of this. The other point of interest in this work is the fact that an absolute floor value is used. The idea of this threshold is to avoid excessive false alarms in dark regions of the imagery.

5.5 Segmenting detectors

In the papers, [70] and [115], Lombardo and Sciotti analyse the performance problems of standard CFAR techniques for ship detection in satellite SAR imagery and describe a new segmentation based technique which addresses these problems. The first paper is a journal paper and provides a more detailed account of the contents of the second paper. These papers represent a substantial and important piece of work and warrant a detailed description here.

Lombardo and Sciotti note that SAR imagery of the ocean is typically inhomogeneous with differences in ocean backscattering behaviour caused by features such as: transitions between regions with different wind conditions; low wind spiral marks; effects due to land areas; altered backscattering due to ocean bottom effects; and the presence of ship wakes. Such inhomogeneities give rise to a large number of false alarms when standard CFAR detectors are used. Hence they propose preceding the CFAR detection stage with a segmentation stage which partitions the imagery into more uniform areas. Both low-resolution quick-look ERS SAR imagery and high resolution X-SAR/SIR-C SAR imagery are considered. Ground truth was obtained by expert humans with the help of accurate land masks.

Lombardo and Sciotti provide a theoretical analysis of three standard CFAR detectors. The standard detectors they analyse are the cell-averaging CFAR detector (CA), the cell average-greatest of (CA-GO) and the cell average-smallest of (CA-SO). The CA detector is as described in (14). The CA-GO and CA-SO detectors are related to the CA detector but are specifically designed to cope with inhomogeneities in the form of step edges. They do this by dividing the background window into two halves and estimating the background mean μ_b in each half. The CA-GO detector uses the maximum of these two values while the CA-SO detector uses the minimum. The background window dividing line can have any specified orientation. Lombardo and Sciotti use a single pixel target window, a 3x3 pixel guard window and a 7x7 background window. It is likely but not explicitly stated that these window sizes are used in the practical applications later in the paper.

Lombardo and Sciotti's theoretical analysis of the performance (PFA and PD) of the CFAR detectors on L-look SAR imagery is based on modelling the background statistics with a Gamma pdf with order parameter L and the ship target statistics (after some derivation) with a Gamma pdf with order parameter $L \times \nu_T$. They analyse detection performance in inhomogeneous areas by considering what happens as the detection window moves across the boundary between two different homogeneous regions. In particular, they consider the case where the background window contains a step edge so that the background is divided into two areas with differing mean intensity values. The calculation of PFA and PD for the CFAR detectors is non-trivial and extensive calculations are needed. The results are summarised in plots of PFA and PD against proximity to the step edge and show that as the transition between regions occurs all three detectors have trouble maintaining both an acceptable false alarm rate and a good probability of detection. Lombardo and Sciotti note that the explicit dependence of PFA and PD on the (effective) number of looks L in their formulae allows detection thresholds to be set and performance studied when operating with SAR images of different quality and resolution.

Lombardo and Sciotti confirm their theoretical findings with evidence from real im-

agery. The CA CFAR detector is seen to have good performance on homogeneous imagery but badly degraded performance on inhomogeneous imagery. While the CA-GO improves on the CA for inhomogeneous imagery, the gain in performance is not significant. When the detection threshold is set low enough to achieve a good PD, too many false alarms are generated in inhomogeneous areas. Lombardo and Sciotti have attempted to reduce the FAR by varying the size and shape of the background window and by taking into account the detected target's dimensions but without success. They note that while previous detection techniques have taken the variance of the background into account and increased the detection threshold in high-contrast areas this is done at the expense of decreased PD.

The problems mentioned above motivate Lombardo and Sciotti to consider prior segmentation as was done in previous target detection work by McConnell and Oliver, [81]. If the step edges and other such inhomogeneities are correctly identified by segmentation and if the CFAR detectors are altered to only use background pixels from the same homogeneous region as the test pixel then the problems of high false alarm rates are alleviated. Interestingly, Lombardo and Sciotti's experience has shown that different algorithms are appropriate for different types of imagery. We discuss the details next.

For low-resolution SAR images with a high effective number of looks, Lombardo and Sciotti use the following segmentation and detection procedure. First, a global threshold is used to split the image into low wind regions (LWR) and normal wind regions (NWR). Next, single pixels and small regions are removed by median filtering. The segmentation is then further refined by identifying "transition regions" (TR) between the two classes and merging them with the low wind regions since their inhomogeneities need to be dealt with in the same way. This last step involves the use of the "normalised-log measure" to take account of texture in the image but is not clearly explained in the paper. The final segmentation result is two classes - NWR and LWR+TR. Different detectors are applied in each class. A CA-CFAR detector is applied in the NWR but modified so that only NWR pixels are used to estimate the background mean. In contrast, a fixed threshold is used in the LWR+TR regions. This is done to avoid problems arising from the high variability of the data and low mean intensity in these regions. The fixed threshold is set as a function of the mean intensity in the neighbouring NWR but the gain factor used (the design parameter t in (14)) is lower than that used in NWR. While the motivations for this algorithm are well explained in the paper, the technical details such as window sizes and gain factors are lacking.

In the higher resolution X-SAR/SIR-C imagery, an improved segmentation technique is needed. Lombardo and Sciotti note that while ship targets occupy more pixels in higher resolution imagery and are therefore easier to detect, the imagery is more inhomogeneous and consequently the false alarm rate is also higher. Lombardo and Sciotti use a simulated annealing segmentation technique, see [24] and [88], which handles speckle well. Following segmentation, two thresholds are applied to the average intensities of the segments so that they are partitioned into 3 classes. These classes have similar texture (as measured by a "normalised log" technique) but differing means. Lombardo and Sciotti state that these classes are adequate for the subsequent detection stage but give no explanation as to why the segments were merged in this way. They also do not specify the values of the thresholds which were used. The resulting segmentation class regions contain no uniform dark "LWR" type regions so they do not use a fixed detection threshold. Instead they use the CA-CFAR in all class regions. Lombardo and Sciotti note it makes sense to use

detectors which take account of the multi-pixel extent of the targets but for various reasons adhere to a single pixel target window. Finally, since ship targets may occupy many pixels, all small "ship sized" class regions are tested against a CA-CFAR threshold derived from the adjacent larger class regions. Precise details of how the latter is done are not given.

Setting the detector design parameter t in (14) to achieve a given PFA is more complicated in higher resolution imagery. Lombardo and Sciotti observe that applying the CA-CFAR as-is gives a FAR 2 orders of magnitude above the design FAR confirming the highly non-Gaussian nature of the region statistics. Hence, they model the region statistics as K-distributed. They derive a new approximate expression for PFA assuming a CA-CFAR is used and the background clutter is K-distributed. This expression allows the correct design parameter for a given false alarm rate to be to chosen. The relationship depends on the K-distribution order parameter ν. In applications to high resolution X-SAR/SIR-C imagery Lombardo and Sciotti note that ν does not vary much within each segmentation class. Consequently, ν is estimated for each class and the design parameter t is set to achieve the desired FAR. A CA-CFAR is then applied in each class region. We note here that in effect, Lombardo and Sciotti are taking the background variance into account and altering the detection threshold appropriately.

Lombardo and Sciotti assess the first algorithm (for low resolution imagery) by applying it to real ERS imagery which was ground-truthed by eye and compare its performance to the standard CA-CFAR and CA-GO CFAR detectors and also to Eldhuset's detector (no reference given but presumably [32]) and OISD, [25]. They report substantially better performance. Similarly they assess their second algorithm (for high resolution imagery) by applying it to X-SAR/SIR-C imagery. Here, they compare performance with the standard CA-CFAR and their first algorithm and show that their second algorithm is substantially better. Thus Lombardo and Sciotti results confirm the usefulness of a prior segmentation stage which allows one to properly estimate both the mean and the order parameter of the background clutter.

Follow up work is presented in [117]. There, Sciotti *et al.* consider the best way of exploiting polarimetric data. They consider three different detection strategies: a) apply single channel detectors to each polarimetric channel and then combine the detections; b) combine the polarimetric channels and then apply a single channel detector; c) apply a multi-channel detection algorithm. Among the single channel detectors used are cell-averaging CFAR (14) and the two parameter CFAR (12). The main thrust of the paper is to demonstrate that non-homogeneous ocean features cause high false alarm rates and consequently an initial segmentation stage improves the performance of the detectors. They observe that the false alarms are most prevalent in non-homogeneous areas of the imagery which contain: (i) very dark segments (ii) edges between areas with different mean back-scattering. In conjunction with the three different detection strategies they consider three different segmentation strategies: 1) segment the single channels separately; 2) segment the combined polarimetric channels; 3) segment the polarimetric channels jointly. The segmentation algorithm used is the algorithm for high resolution imagery in [70] as just discussed. They assess the performance of the various detection schemes by applying them to several different L-band SIR-C SAR images. Their results clearly demonstrate a large performance improvement when the segmentation stage is introduced.

5.6 Likelihood ratio detectors

So far, the detectors we have discussed have been based on statistically modelling the background and looking for pixels with a low probability of belonging to the background. However, for an optimal detector, it is necessary to take the target into account too. Without going into the details, such considerations lead to the likelihood ratio test or Neyman-Pearson detector

$$\frac{p(x|T)}{p(x|B)} > \lambda \qquad \Longleftrightarrow \qquad \text{target}$$

where x is the data under test, $p(x|T)$ is the pdf of x given that a target is present, $p(x|B)$ is the pdf of x given that only background is present and λ is a fixed threshold. Generally, λ is chosen to give a specified probability of false alarm.

The likelihood ratio detector requires an explicit knowledge of the pdfs $p(x|T)$ and $p(x|B)$. However, in target detection work, as we have seen above, these pdfs vary from location to location and are not known explicitly. The usual approach is to model them parametrically and to estimate the parameters from the data. Thus we assume parametric models $p(x|T, \theta_T)$ and $p(x|B, \theta_B)$ where θ_T and θ_B are the vectors of pdf parameters. When θ_T and θ_B are estimated using maximum likelihood methods, we obtain the generalised likelihood ratio test detector

$$\frac{\max_{\theta_T} p(x|T, \theta_T)}{\max_{\theta_B} p(x|B, \theta_B)} > \lambda \qquad \Longleftrightarrow \qquad \text{target} \qquad (26)$$

where again λ is chosen to give a specified probability of false alarm.

Generalised likelihood ratio detectors are not commonly used in target detection work because targets are hard to model well statistically. Ship targets, for instance, come in many sizes, shapes and orientations and for this and other reasons their appearance in SAR imagery is highly variable. However, Brizi et al., [10], have proposed a simple but useful model. The associated GLR detector has been considered for ship detection by Sciotti et al., [116] and [117]. We describe it next.

Brizi et al., [10], consider single look complex (SLC) SAR imagery. They use a target window of N pixels and arrange the complex pixel values in a column vector $\mathbf{x}$. Under the hypothesis (T) that the window is centred on a target, $\mathbf{x}$ is assumed to contain target echos plus thermal noise and under the alternative hypothesis (B), that the window is not centred on a target, $\mathbf{x}$ is assumed to contain clutter echos plus thermal noise. The model which Brizi et al. propose is that the target, clutter and thermal noise samples are all extracted from zero mean Gaussian random processes with variances σ_t^2, σ_c^2 and σ_n^2, respectively. This is a commonly used model for clutter samples from SLC SAR imagery, at least. Brizi et al. also assume that the samples are independent from pixel to pixel which is certainly reasonable for low resolution imagery. To aid in modelling the background clutter, Brizi et al. also make use of K secondary data $\mathbf{y}_1, \dots, \mathbf{y}_K$, each of dimension N, which are assumed to be target free. The secondary data are arranged into a matrix $\mathbf{Y} = [\mathbf{y}_1, \dots, \mathbf{y}_K]$. Thus the generalised likelihood ratio in (26) takes the form

$$\frac{\max_{\sigma_t^2, \sigma_c^2, \sigma_n^2} p(\mathbf{x}, \mathbf{Y}|T, \sigma_t^2, \sigma_c^2, \sigma_n^2)}{\max_{\sigma_c^2, \sigma_n^2} p(\mathbf{x}, \mathbf{Y}|B, \sigma_c^2, \sigma_n^2)} = \frac{\max_{\sigma_t^2, \sigma_c^2, \sigma_n^2} \left\{ p(\mathbf{x}|T, \sigma_t^2, \sigma_n^2) \cdot p(\mathbf{Y}|\sigma_c^2, \sigma_n^2) \right\}}{\max_{\sigma_c^2, \sigma_n^2} \left\{ p(\mathbf{x}|B, \sigma_c^2, \sigma_n^2) \cdot p(\mathbf{Y}|\sigma_c^2, \sigma_n^2) \right\}} \qquad (27)$$

Substituting the pdfs in and performing the maximisations leads to the detector

$$\frac{\left(\sum_{k=1}^{K}\sum_{n=1}^{N}|y_{k,n}|^2 + \sum_{n=1}^{N}|x_n|^2\right)^{K+1}}{\left(\sum_{k=1}^{K}\sum_{n=1}^{N}|y_{k,n}|^2\right)^{K}\left(\sum_{n=1}^{N}|x_n|^2\right)} > \lambda \qquad \Longleftrightarrow \qquad \text{target} \qquad (28)$$

which Brizi *et al.* refer to as the Gaussian Generalised Likelihood Ratio Test (G-GLRT). Brizi *et al.* use $K = 3$ but make no comment on how the secondary data $\mathbf{Y} = [\mathbf{y}_1, \ldots, \mathbf{y}_3]$ are obtained. It is of interest to note here that Brizi *et al.* also extend the target model to take account of the information available in the target radar shadow. Thus their extended detector looks for bright target pixels next to dark target shadow pixels. However, moving ships are shifted in azimuth from their shadows and therefore this extension is not relevant to ship detection.

The G-GLRT detector of Brizi *et al.* was developed into a typical CFAR sliding window approach by Lombardo *et al.* in [71]. They use a 5x5 pixel target window, a 11x11 pixel guard window and a 15x15 pixel background window. Thus the secondary data are no longer of the form specified by Brizi *et al.* and instead are formed from the pixels in the annulus between the background and guard windows. This change requires the G-GLRT detector to be reformulated. Lombardo *et al.* do this for the case where the target shadow is modelled as well as the target itself. Lombardo *et al.* also work with intensity data rather than complex data. By extrapolating from the reformulation given by Lombardo *et al.* in [71] we believe that the G-GLRT detector (28) becomes

$$\frac{(N_b\mu_b + N_t\mu_t)^{N_b+N_t}}{\mu_b^{N_b}\mu_t^{N_t}} > \lambda \qquad \Longleftrightarrow \qquad \text{target} \qquad (29)$$

where μ_t and μ_b are the mean intensity (power) in the target and background windows, respectively, and N_t and N_b are the number of pixels in the target and background windows, respectively.

The G-GLRT detector has been considered for ship detection by Sciotti *et al.* in [116] and [117]. While their main aim in these papers has been to explore the best method of exploiting polarimetric data, they still compare the performance of various detectors on single channel imagery and fused multi-channel imagery. The results in [116] with simulated data show that the G-GLRT detector performs better than the cell averaging CFAR detector (16) (which they refer to as the Power Ratio detector). However, the best detection performance is given by the PG-GLRT detector which is designed to exploit polarimetric information. We will discuss the PG-GLRT detector later in this report in the section on polarimetric data. In [117], performance is assessed on SLC quad-pol L-band SIR-C SAR data. Detection thresholds are set at the extreme values which still ensure the detection of all visible ship targets. Performance is then assessed by comparing the false alarm rates. While their results for polarimetric data which has been fused into a single image verify that a G-GLRT detector out-performs a cell averaging CFAR detector, they are both out performed by a two parameter CFAR detector (15). However, their results also show that if an initial segmentation stage is included so that the detectors are applied to homogeneous regions only then the G-GLRT detector can out-perform the two parameter CFAR detector.

5.7 The SUMO detector

The fisheries control group at the Joint Research Centre (JRC) of the European Community are developing a software system for the automatic detection of ships in satellite SAR imagery. The system is known as Search for Unidentified Maritime Objects (SUMO). Its development has been documented by Schwartz *et al.* in [114] and Kourti *et al.* in [61] and [60]. While the paper by Schwartz *et al.* describes two algorithms for discriminating false alarms from true detections and therefore more properly belongs in the section on discrimination algorithms, we review it here since one of the algorithms is very much tied to the detection algorithm used. It also contains further details of the detection algorithm.

The work of Kourti *et al.*, [61], was aimed at generally assessing the utility of Satellite SAR for fisheries control, it also includes the details of an algorithm for automatic processing of the imagery. The study revealed that of the available Radarsat modes, ScanSAR narrow imagery is the best for fisheries monitoring since it covers a wide area with a sufficient resolution. Consequently, a total of 60 ScanSAR narrow images from 3 different fisheries were analysed. Groundtruth was provided by the Vessel Monitoring System (VMS). As an interesting aside, we note the comment in [61] that JRC's own software package, G-PROC, (G-PROC was developed specifically for SAR analysis in the early 1990s) was awkward and unreliable, therefore the ENVI/IDL package was chosen.

The SUMO ship detector is a type of template matching. The template is designed to match the appearance of ships in SAR imagery. It consists of a 4x4 pixel window with three different thresholds as described in step e) below and displayed graphically in Figure 2. The detection algorithm consists of the following steps. a) Mask out the land. b) Partition the image into tiles of size 2500 lines by 400 pixels; c) Calculate the mean μ and standard deviation σ of the pixel values in the tile being processed. d) Calculate thresholds for the search template by

$$T1 = 3.0 * (\mu + \sigma) \qquad T2 = 1.9 * (\mu + \sigma) \qquad T3 = 1.4 * (\mu + \sigma) \qquad (30)$$

The threshold factors 3.0, 1.9 and 1.4 were found empirically. They were region dependent and the combination 2.5, 1.3 and 1.1 was better in areas with low ocean clutter levels. Work on automatically calculating the optimal values for a particular area is underway. e) A detection is recorded at pixel (i, j) if:

- The pixel values at (i, j), $(i+1, j)$, $(i, j+1)$ and $(i+1, j+1)$ are above $T1$,

- and the values at $(i-1, j)$, $(i-1, j+1)$, $(i+2, j)$, $(i+2, j+1)$, $(i, j-1)$, $(i, j+2)$, $(i+1, j-1)$ and $(i+1, j+2)$ are above $T2$,

- and the values at $(i-1, j-1)$, $(i-1, j+2)$, $(i+2, j-1)$, $(i+2, j+2)$ are above $T3$.

f) Apply the window at every pixel (although there is some ambiguity in the wording of this step). g) Group neighbouring detected pixels into one. h) Convert detected pixel coordinates into geographic location so validation against the groundtruth is enabled. i) Remove detections with more than five pixels in a row or column above $T1$. Fishing vessels are not longer than 150m and therefore are unlikely to generate such detections. Improvements on this classification module are underway. j) Apply corrections to georeferencing

T3	T2	T2	T3
T2	T1	T1	T2
T2	T1	T1	T2
T3	T2	T2	T3

Figure 2: The SUMO detector template

using better satellite orbit parameters or known ground points when and where possible. This revealed consistent geolocation errors of about 1 nautical mile in one fishery.

This detector is hard to explain in rigorous theoretical terms. While its use of thresholds based on the mean and standard deviation of the surrounding ocean is similar to that of the two parameter CFAR detector (12) and hence understandable, and while the extension of this idea to a template over which different thresholds are applied is also fairly natural (given the desire to match the appearance of ship in SAR imagery), the calculation of the thresholds themselves is unusual and hard to explain. We believe a reporting mistake has been made and that (30) should be replaced by

$$T1 = \mu + 3.0 * \sigma \qquad T2 = \mu + 1.9 * \sigma \qquad T3 = \mu + 1.4 * \sigma \qquad (31)$$

so that the thresholds are calculated according to the two parameter CFAR detector formulae in (12). Corroborating evidence that this is the case was found in a presentation on the JRC web site, [57]. Note that the threshold factors 3.0, 1.9 and 1.4 used here are much lower than the usual value of 5.0 in a two parameter CFAR detector. While the low thresholds have the potential to produce a high false alarm rate, the use of a template which takes adjacent pixels into account will have the opposite effect. However, we warn here that using a template in this way assumes the ship will have multi-pixel extent and so is not appropriate when looking for small ships in low resolution imagery.

In the results section of [61], Kourti *et al.* note that apart from the detection algorithm, their exercise is similar to that of Vachon *et al.*, [134]. The results of Kourti *et al.* clearly indicated that wooden and fibreglass vessels do not show up well in SAR imagery. Several other detection problems are mentioned: weather effects (high wind speed generates noisy backgrounds creating detection problems); incidence angle effect (the near swath ocean clutter is brighter than the far swath clutter); and image errors (bright spots at the edges of different beam modes). We report their conclusions by quoting from the paper: *Results from different studies confirm that fishing vessels longer than 26 m length have a 92% probability of being detected on ScanSAR imagery. ScanSAR imagery offers a good compromise between resolution and covered area. The comparison of the detected vessels positions with VMS position reports match well. More than 73% of the vessels could be unambiguously identified in the Flemish Cap study and 92% in the North Sea. Only in the Azores, where fibreglass and wooden vessels predominate, was it difficult to detect the vessels subject to VMS. The mean distance between detected position and VMS position*

T3	T2	T2	T2	T3
T2	T1	T1	T1	T2
T3	T2	T2	T2	T3

Figure 3: A second SUMO detector template

was about 0.3 nautical miles. Other observations of interest in the paper are: trawling vessels have larger wakes; azimuthal travelling vessels are brighter than range travelling ones.

Further information on the SUMO detector is given in the follow up paper by Schwartz *et al.*, [114]. Schwartz *et al.* explain that near real-time processing is a critical requirement of SUMO and hence a very simple detector was chosen. In addition to the 4x4 pixel template described above, a second 3x5 template is presented, see Figure 3. For this template, step e) of the algorithm becomes: e) A detection is recorded at pixel (i, j) if:

- The pixel values at $(i, j-1)$, (i, j), $(i, j+1)$ are above $T1$,

- and the values at $(i, j-2)$, $(i-1, j-1)$, $(i-1, j+1)$, $(i, j-1)$, $(i, j+1)$, $(i+1, j-1)$, $(i+1, j+1)$ and $(i+2, j+2)$ are above $T2$,

- and the values at $(i-1, j-2)$, $(i+1, j-2)$, $(i-1, j+2)$, $(i+1, j+2)$ are above $T3$.

Unfortunately, there is no discussion of how the two templates are used and what the difference between them is — presumably both are run over the imagery and their detections are combined. We note here that the 3x5 template is the only reference in the ship detection literature to a non-square detector window.

The main thrust of the paper by Schwartz *et al.* is the study of algorithms for identifying and removing the two main types of false alarms produced by the SUMO detector. Their analysis revealed that most false alarms arise from two circumstances:

- Large objects, like oil rigs, which for certain aspect angles have bright ship-like sidelobe patterns

- Very local wind turbulence effects which can produce bright ship-like noise patterns.

Schwartz *et al.* note that false alarms in the first category could easily be eliminated if the locations of the large objects are mapped and the SAR image registered with the map. However, that is not possible owing to the lack of precise satellite orbit parameters. Schwartz *et al.* also comment that modelling and predicting the sidelobe patterns is not an option. Hence an *ad hoc* method is employed. It involves simply grouping together any detected pixels which are less than 8 pixels apart. In this way the false alarms due to sidelobes are incorporated into the main detection. The distance of 8 pixels was selected heuristically. It was considered suitable for the detection problem at hand since ships in

the open sea are usually more than 300 m apart which corresponds to a distance of 12 pixels.

For the second type of false alarm, Schwartz *et al.* note that local wind turbulence effects are normally contained within a 100x100 pixel area. Three characteristics of such effects are listed. The last of these was found to be best for discrimination purposes. It states that the number of local maxima close to the threshold value $T2$ is large. Therefore, a 100x100 pixel image chip is taken from around each detection and the number of local maxima above $T2$ is counted. If that number is larger than 750, the detection is rejected as a false alarm. Test data was used in setting the threshold at 750. Schwartz *et al.* report that this algorithm separates false alarms from true detection well and provides quantitative performance data.

In interesting related work, Aresu and Schwartz, [4], have considered the effect of the SAR processor on ship detection and ship RCS. Amongst other things, they note that while logarithmic output scaling is useful for oil spill detection, it is not suitable for ship detection. JRC researchers have confirmed that SUMO does not work well on logarithmically scaled SAR imagery and that linear or square root scaling is preferable. This point was made in a presentation to the EU Fisheries Expert Group meeting in June 2003 which is available from the JRC web site, [57]. We note that this result is not really surprising. It is commonly assumed that the gamma distribution with an order parameter equal to the effective number of looks is a reasonable model for intensity SAR imagery. Hence the two parameter CFAR detector, and consequently SUMO which is based on it, will be approximately CFAR on intensity SAR imagery. Taking logarithms of the imagery will change the statistics and it is unlikely that the detectors will remain CFAR.

5.8   The Lockheed Martin Canada detector

An interesting prescreening algorithm has been proposed by researchers at Lockheed Martin Canada. It is documented in the paper by Gagnon *et al.* [41]. While there is no account of using this prescreener for the detection of ships in SAR imagery, it is clear that that is one of the intentions. Coastal surveillance is cited as one of the motivations and the paper includes the description of a system for ship recognition in high resolution SAR imagery. One of main the attractions of the prescreener is that it is based on a multiresolution analysis and hence is especially applicable to high resolution imagery. It is also reported to be fast. The research was undertaken in collaboration with Universities (Montreal and Queens) and the Canadian Department of National Defence.

Gagnon *et al.* [41] state that the prescreener is a practical multiresolution algorithm for target detection in uniform clutter and that it is based on the assumption that man-made objects are easily detectable at low resolution because their scattering is more persistent than that of natural objects. Other more sophisticated prescreeners were considered but the adopted one was chosen as a good tradeoff between algorithm complexity and low FAR.

The algorithm is applied to stripmap SAR imagery and the steps involved are: application of a discrete wavelet transform (DWT); extraction of the relevant subbands; formation of low resolution amplitude images; application of a Rayleigh CFAR detector

to the amplitude imagery; extraction of corresponding regions of interest (ROI) in the original imagery; delimitation of the ROIs; clustering of the ROIs; and application of a Weibull or K distribution CFAR detector to the final ROIs. We elaborate on these next.

The DWT uses symmetric real filters which behave like bi-orthogonal wavelets. Unfortunately no details are given on the actual wavelets used. As usual, at each resolution level there are 4 wavelet subbands — LL, LH, HL and HH — where L denotes the low pass filter and H the high pass filter. The operator then chooses which resolution levels are to be subjected to further analysis. This choice depends on the size of the targets being sort. For each chosen resolution, an amplitude image I is formed by setting

$$I = \sqrt{LH^2 + HL^2} \tag{32}$$

If the resolution is chosen correctly then the targets of interest should generate bright single pixels in I. It is assumed that the wavelet cofficients in LH and HL are Gaussian distributed with the same variance. It follows that I must be Rayleigh distributed. Gagnon notes that this assumption is roughly true in practice provided the clutter is homogeneous. Consequently, a Rayleigh CFAR detector is applied to I.

According to Ravid $et\ al.$ [99] if the Rayleigh distribution is written as

$$f(x) = \frac{2x}{B^2} \exp\left(-\frac{x^2}{B^2}\right) \tag{33}$$

then the corresponding CFAR detector is

$$x_t > Bt \qquad \text{where} \qquad B = \sqrt{\frac{1}{M}\sum_{i=1}^{M} x_i^2} \tag{34}$$

where $x_1, x_2, \ldots, x_M$ are the samples from the background clutter used to estimate the parameter B. Further, the PFA for this detector is given by

$$PFA = \left(1 + \frac{t^2}{M}\right)^{-M}. \tag{35}$$

We note here that for Rayleigh distributed clutter, both the mean and standard deviation are proportional to B and hence this detector will tend to behave like a cell-averaging CFAR detector and also a two parameter CFAR detector. Gagnon $et\ al.$ apply this detector to I with the standard moving window setup shown in Figure 1. It is implicit that the target window is a single pixel. No sizes are given for the guard and background windows but it is stated that they are square and small since target dimensions are small at low resolutions.

While a simple application of the Rayleigh CFAR detector to I with an appropriate PFA will identify the locations of potential targets of interest, Gagnon $et\ al.$ add a refinement in order to extract well-placed regions of interest (ROI). The exact details of how this is done are not clear from the paper but our interpretation is as follows. The Rayleigh CFAR detector is applied to I twice. The first pass has a high PFA and identifies only the very bright pixels. These pixels are assumed to be target pixels and they are assigned a square ROI in the original image. The size of the ROI is an (unspecified)

function of the size of the targets being sought. The second pass is aimed at detecting nearby pixels in I which are not as bright as those detected in the first pass but still likely to be associated with them. Thus the second pass uses a lower PFA and only searches around the pixels detected in the first pass. The position of the ROIs in the original image is then adjusted by weighting them according to the intensity of the pixels detected in the second pass.

Next, any overlapping ROIs in the orginal image are clustered into one larger ROI under the assumption that they correspond to a large target which is not point-like at the selected wavelet subband resolution. Then, in the final step of the algorithm, another CFAR detector is applied to the ROIs in order to eliminate false alarms. It is noted that this CFAR detector should be appropriate to the statistics of the orginal image. It is suggested that for ocean clutter a K-distribution based CFAR detector is appropriate while for land clutter a Weibull CFAR detector is better. No mention is made as to how the problem of large targets in high resolution imagery is dealt with in this last step — single pixel detectors are obviously not appropriate.

We refer to the complete prescreening algorithm as the LM detector. Gagnon *et al.* state that LM detector with its combination of detecting ROI centres in the wavelet subbands followed by CFAR detection in the original image, results in a much lower FAR than the usual method of only applying CFAR detection in the original image. A graphic illustration of this point is given with the detection results of the two methods being displayed on a SAR image of tanks and trucks on land. Unfortunately detailed performance results for the LM detector are only presented for land imagery. Further, apart from the illustration just mentioned, no comparisons with other detectors are made. However, it is presumed that more extensive comparisons with simple CFAR detection in the whole image have been made (for land imagery at least) and that the LM detector is superior. We conclude with the remark that the imagery used had a resolution of one-foot, highlighting the applicability of the technique to high reolution imagery.

5.9   DSTO detectors

A study of algorithms for the detection of small boats in Radarsat imagery has been previously conducted at DSTO by Hunt *et al.* [51]. That study was aimed at assessing the performance of three stochastic model-based detectors which were developed earlier at DSTO [20]. These detectors are known as: the "conditional probability test" detector; the "significance test" detector; and the "conditional probability significance test" detector. Their performance was compared with the two parameter CFAR detector. The theory underpinning the stochastic model-based detectors assumes the imagery has Gaussian statistics. Consequently, a preprocessing technique which transforms the histogram of the whole image so that it becomes Gaussian was also considered. We note that the DSTO detectors and the preprocessing transformation are much more computationally expensive than the two parameter CFAR detector. They also assume that the background statistics are stationary over the image.

The test data consisted of part of a single Radarsat Fine mode image of Darwin harbour with a resolution of 9–11 m in range and 9 m in azimuth. The image area under analysis contained 15 well ground-truthed fishing boats. Eleven of the boats were wooden

and the rest were wood and steel. The boat lengths varied between 8 and 15 m. It was noted that the choice of detector threshold and detector window sizes had a considerable impact on performance. Consequently, a range of detector thresholds and window sizes were considered. We note here however that the two parameter CFAR target window was not varied and remained a single pixel throughout.

The authors report that the stochastic model-based detectors produced substantially better results than the two parameter CFAR detector and that even better results were obtained by applying the preprocessing technique. The "significance test" and the "conditional probability significance test" were particulary recommended. They achieved a PD of 93% with zero FAR when used in conjunction with the preprocessing step. However, the authors note that the test data is limited in size and are reluctant to draw too many conclusions from their results. We quote from the executive summary: *Since, however, only a single image with a relatively small number of targets was available for analysis, the results here are only generally indicative of the likely performance of detection routines in a real system.* We add here that the boats were moored in a harbour and hence the scenario was not representative of an open ocean environment. The authors recommend further testing with more ground-truthed targets and a variety of weather conditons. They also comment on the need for coastline detection in order to discount false alarms arising from detections on land.

5.10 Other detectors

For the sake of completeness, we discuss in this section other detectors which have been used for ship detection. The most convincing of these is a brief reference by Yeremy *et al.* in [143] to the use of the fourth moment, the normalised kurtosis, $a_4 = m_4/m_2^2$ where m_i is the i-th moment of the intensity. It was found to be suitable for ship detection in open ocean conditions and successfully detected 67 ships out of a possible 70. It is noted that the method is not computationally intensive or sophisticated. However, no technical details are given in [143] as to how the normalised kurtosis is used and the companion paper referred to there appears to have not been published.

In a pair of related papers, Benelli *et al.*, [8], and Argenti *et al*, [5] have developed a ship detection system with several stages. The proposed processing chain involves a parallel search for both ships (bright spots) and ships wakes (linear structures) followed by a cross-validation of the wakes against the ships. The system is based on fuzzy decisions and these are used to assign confidence values to the detections. An initial speckle filtering stage is included. Many different speckle filters were tested. Ships are detected and delineated at the same time by a region growing technique which is seeded with the very brightest pixels. The number of ship-seeds to start with is taken as the maximum number of ships expected times the expected average size. The ship-seeds are processed in order of brightness. Neighbouring bright pixels are added to the ship-seeds one at a time. The bright pixels which are eligible to be added are identified by taking a p-tile of the right tail of the image histogram. While the method is at heart simple, the exact details are not clearly spelled out in the paper. In particular, it is stated that the value of p is adaptable and is computed on the basis of the image histogram but no details are given. The output of this process is a list of potential "ship-blobs". The next step is aimed at identifying

false alarms among the ship-blobs. It uses a classifier based on fuzzy set theory. The classification features are area, mean grey level and elongation. The classifier is designed so that true ships have a high fuzzy index and false ships have a low fuzzy index. Further false alarm reduction occurs in the final stage of the system where ships and wakes are paired. The pairing is based on minimising the distance between the ship centroid and an extreme of its associated wake. The number of ships can differ from the number of wakes. Following pairing a coupling coefficient is calculated for each ship-wake pair. This coefficient when multiplied by the fuzzy index of the ship provides a final confidence measure for each detected ship-wake pair.

In the SACLANT report [7], Askari and Zerr investigate two new ship detectors. The first is referred to as the Neural-Network-Dempster-Shafer technique. It involves sub-dividing the image into sub-blocks and extracting four separate 1-D profiles with different orientations from each sub-block. The profiles are processed by a neural-network which measures (in terms of belief) the extent to which each profile matches a pre-defined Gaussian shape. The individual beliefs are fused using Dempster-Shafer rules to derive a final probability of the sub-block being a target. The second detector, referred to as the MM algorithm, operates on a binary image produced by setting all pixel values below 255 to zero. It employs an erosion algorithm and then groups the detected pixels into clusters. Radarsat imagery was used to test the algorithms. Curiously, one of Askari and Zerr's findings was that Radarsat standard beam S6 was better for ship detection than Radarsat ScanSAR and they attribute this to the better radiometric resolution of the 16-bit product over the 8-bit product.

Osman and Blostein, [91], use an image segmentation approach to ship detection in airborne SAR imagery. Their segmentation algorithm is based on a neural clustering scheme referred to as "probabilistic winner-take-all" (PWTA). Once the image has been partitioned into segments, ship segments are identified by assessing the similarity between segments. It is assumed that there is only one ship segment and it is hypothesised that the similarity between ocean segments is much higher than the similarity between ocean segments and the ship segment. Similarities are measured with the Bhattacharyya distance. The ship segment is identified as the one with the largest average Bhattacharyya distance B_s from the other segments. A special problem in this approach is determining the number of segments K to use. The solution offered is to try a range of values for K and to choose the one which maximises B_s. The PWTA algorithm was shown to perform better than other segmentation techniques and an average ship detection accuracy of 87.5% was reported.

Leung *et al.*, [66], use a radial basis function (RBF) neural network for ship detection. They justify the use of a neural network on the basis that other studies have shown ocean clutter to be chaotic rather than purely random. The essential idea is to model the surrounding background ocean clutter with the RBF and then predict the current pixel value with it — poorly predicted pixels are anomalies and therefore likely ships. They use a genetic algorithm to fine tune some of the RBF parameters. A CFAR threshold is set for the prediction error by using histograms to estimate the prediction error pdf. The technique is applied to Radarsat SAR imagery. The results are reported to be reliable although no comparison with standard detectors is made.

A novel approach to target detection in SAR data is suggested in the paper [75] where

a ship example is used. Unfortunately the paper is poorly written and is therefore hard to assess. The essential idea is to detect targets in the so called "range-Doppler" domain before azimuth compression of the data. The technique involves target detection in the 1-D range profiles followed by target confirmation based on accumulation of azimuth information.

Another brief and poorly written paper is the one by Wang *et al.*, [139]. This paper deals with very high resolution SAR imagery. Further, the technique described appears to be more suitable for target delineation and classification rather than detection. A related follow up paper is [21]. Similar comments apply to it.

6   Prescreening — polarimetric SAR

Most of the development of ship detection algorithms and systems to date has concentrated on single channel SAR imagery. This is probably due to the fact that most researchers work with data from commercial satellites such as Radarsat-1 and ERS-1/2 and those systems only collect single channel imagery. However, the recently launched Envisat and the soon to be launched Radarsat-2 will provide dual and quad-polarised SAR data, respectively. It is anticipated, see for instance [142], that the additional information contained in such polarimetric data will improve ship detection (as well as classification and even identification) capabilities. In fact, in [142], Yeremy *et al.* state that current single channel ship detection algorithms have high false alarm rates and require human supervision to achieve acceptable results but go on to suggest that with polarimetric data, ship detection algorithms will be developed which can reduce false alarm rates to an acceptable level by correctly classifying the false alarms. (As an aside we note here that in the report, [145], Yeremy *et al.* consider the military applications of polarimetric interferometric SAR (PolInSAR) methods where even more independent information is available, however, despite a reference to ships in the abstract no details were found.) Consequently, there has been much recent research on ship detection with polarimetric SAR and detection algorithms are still being developed. We report on such research in this section. We begin with a very brief overview of polarimetric SAR.

The polarisation of a radar wave is described in terms of its electric field vector and more specifically, it refers to the locus of that vector in the plane perpendicular to the direction of propagation. Thus in order to discuss polarisation properties it is better to consider electric fields rather than radar power. The polarisation state of an arbitrary wave is usually described by expressing its electric field vector E in terms of the standard polarisation basis consisting of horizontal (H) and vertical (V) linear polarisations. Thus $E = E_H \hat{h} + E_V \hat{v}$ where $\hat{h}$ and $\hat{v}$ are unit vectors in the "horizontal" and "vertical" directions. If we let $E^i = E_H^i \hat{h} + E_V^i \hat{v}$ denote the incident wave on a scatterer and $E^s = E_H^s \hat{h} + E_V^s \hat{v}$ the scattered wave observed at a distance R from the scatterer then using the superposition leads to:

$$\begin{pmatrix} E_H^s \\ E_V^s \end{pmatrix} = \frac{\exp(i2\pi R/\lambda)}{R} \begin{pmatrix} S_{HH} & S_{HV} \\ S_{VH} & S_{HH} \end{pmatrix} \begin{pmatrix} E_H^i \\ E_V^i \end{pmatrix} \tag{36}$$

where λ is the radar wavelength. The terms S_{HH}, S_{HV}, S_{VH}, and S_{HH} are referred to *complex scattering amplitudes* or *complex reflectivities*. We emphasise that complex

reflectivities are complex numbers and hence incorporate phase as well as amplitude information. It follows from (36) that the *scattering matrix*

$$S = \left[\begin{array}{cc} S_{HH} & S_{HV} \\ S_{VH} & S_{VV} \end{array} \right] \tag{37}$$

contains all the information needed to calculate the complex reflectivity for any pair of transmit/receive polarisation combinations. By taking account of the relationship between the power of a radar wave and its electric field strength it is not hard to deduce that the radar cross-section of a target is given by

$$\sigma_{PQ} = 4\pi |S_{PQ}|^2$$

where S_{PQ} is the complex reflectivity for the combination of transmit and receive polarisations specified by P and Q.

The design of the radar transmit (and receive) antenna determines the transmitted (and received) polarisation. Single channel SAR systems are designed with a single antenna which is used for both transmitting and receiving. In this case, the radar will only produce a single co-polarised channel. Commonly that channel is either S_{HH} or S_{VV}. However, by using a pair on antennas, one oriented horizontally and the other vertically it is possible to design radars which measure all four channels in the scattering matrix (37). Such radars are referred to as fully polarimetric or quadrature polarised (quad-pol). As already noted, fully polarimetric SAR systems allow the calculation of RCS for any pair of transmit/receive polarisation combinations. We should add here that the *polarimetric reciprocity principle* states that $S_{HV} = S_{VH}$. It is generally believed that the reciprocity principle holds for most SAR imaging situations. Consequently it is common to average S_{HV} and S_{VH} in quad-pol SAR data and have only one cross channel.

An important aspect of polarimetric radar theory is the *degree of polarisation*. A electromagnetic wave can be decomposed into two components: a polarised component in which the polarisation is deterministic and an unpolarised component in which the polarisation state changes randomly with time. The ratio of the power between these two components is called the *degree of polarisation*. While radars are designed to transmit fully polarised radiation, the backscattered signal will not necessarily be so. According to (Microwave Remote Sensing (MRS) pg 141), depolarisation is primarily caused by: quasi-specular reflection for corner reflectors; multiple scattering from rough surfaces; and multiple volume scattering due to inhomogeneities.

6.1 Optimal polarisations for ship detection

In this section, we review research efforts comparing different transmit and receive polarisation combinations for ship detection. The most comprehensive results in this respect come from the comparison of ERS SAR data with Radarsat-1 data. ERS data is VV polarised while Radarsat-1 is HH. It is widely accepted that the Radarsat-1 data is better for ship detection (as opposed to ship wake detection), see for instance [110] and [127]. Note however that other factors such as resolution, incidence angle and radar frequency also come into play and care needs to be taken in making such comparisons. The

explanation for Radarsat's superior performance is that HH polarisation has a lower ocean clutter signature than does VV and further, it decreases more rapidly with increasing incidence angle. Hence, ship-sea contrast is higher for HH polarisation than VV and detectability in HH data improves with increasing incidence angle. This also explains why ship wakes and other ocean surface features are rarely seen in HH imagery and why VV imagery is much better for imaging such features.

The most rigorous research effort into the best polarisation combination for ship detection is that reported on by Yeremy *et al.* in [142]. One of the two main aims of the study was to contrast and compare the differences in ship detection potential of the four linear polarisation channels (HH, HV, VH, VV). This was done by measuring and comparing target to clutter ratios (TCR). Yeremy *et al.* define TCR by

$$\text{TCR} = \frac{\sigma_T}{N\sigma_C^0}$$

where σ_T is the total target RCS, N is the number of pixels occupied by the target and σ_C^0 is the ocean clutter's mean RCS. Targets were delineated from the clutter by either supervised delineation or supervised thresholding of the HH (or HV for small incidence angles) channel. A 10 pixel guard band was used when extracting σ_C^0.

The data for the study of Yeremy *et al.*, [142], was collected with the Convair-580 C-band SAR system developed by CCRS. The imagery was SLC with azimuth and slant range sample spacings of 0.42 m and 4 m. Two separate data sets were analysed. The first was collected during the MARCOT'98 and the second during the CRUSADE'00 trials. The MARCOT'98 trail had many ships of opportunity which were visually ground-truthed while the CRUSADE'00 had fewer ships which were rigorously ground-truthed. The data was edited to remove saturation and underflow problems. Further discussion of the CRUSADE'00 trail data is given by Yeremy *et al.* in [143]. It is noted in this later paper that the Convair-580 C-band SAR data have similar characteristics to those planned for Radarsat-2.

The results of the study of Yeremy *et al.* indicate that of the four channels (HH, HV, VH, VV) and for a range of incidence angles from $15 - -70^o$, HH data are best for ship detection when the incidence angle is larger than 45^o while the cross-polarised (HV and VH) data are best when the incidence angle is less than 45^o. Yeremy *et al.* note that this result corrects earlier suggestions (based on a limited range of incidence angles) that HH is always best. Their results also verify that as expected from the reciprocity theorem, the HV and VH data are very similar. As an interesting aside, we repeat their comment that many unknown variables such as ship size, wind speed and direction, effect the TCR results and the degree of variability emphasises how complicated and difficult single channel ship detection is.

The CRUSADE'00 trail data used by Yeremy *et al.*, [142], was also used by Hawkins *et al.*, [45]. Hawkins *et al.* give a more detailed account of the radar setup and data acquisition strategy and a more detailed breakdown of the results but only conduct a preliminary analysis. The conclusions of Yeremy *et al.* regarding the comparison of the four linear polarisation channels (HH, HV, VH, VV) are supported by the Hawkins analysis. While Yeremy *et al.* only discuss RCS, Hawkins *et al.* also consider phase information. The results of Hawkins *et al.* indicate that the relative phase between the different polarimetric channels may be useful for ship detection purposes.

In work related to that of Yeremy *et al.* [142], and Hawkins *et al.*, [45], Touzi *et al.* have also compared different transmit and receive polarisation combinations for ship detection, see [124], [123], [119] and [118]. In the most recent of these papers, [124] and [123], Touzi *et al.* confirm the finding of Yeremy *et al.* that of the three linear polarisations (HH, HV and VV), HV gives highest ship-sea contrast at low incidence angles while HH is better at high incidence angles. However, Touzi *et al.* estimate the crossover incidence angle to be 55^o rather than 45^o. Touzi *et al.* also comment that the effectiveness of HV data seems to be robust to rough sea conditions (at 35 and 45 incidence angle at least.) These results are based on two different data collections. Touzi *et al.* refer to the first as Halifax'98 while the second was CRUSADE'00. It is not clear how reliable these results are due to the limited amount of data presented. We also add that Touzi *et al.* do not define what is meant by ship-sea contrast nor specify how it was measured.

In [123] Touzi *et al.* also look at polarisation channel phase information. They find that the HH-VV channel phase difference looks to be promising for ship enhancement and is certainly better than the HH-HV and VV-VH phase differences. However, they also find that the HV channel is still better than phase information. In the earliest papers, [119] and [118], Touzi discusses the co- and cross-polarisation signatures of ships and open ocean. These discussions, while brief, do provide some insight into the use of polarimetric SAR for ship detection. Finally, we note that Touzi [118] argues that circular polarisations minimise single bounce and maximise double bounce backscattering and hence should be good for ship detection. He states that the ship-sea contrast for circular polarisation looks better than for HH polarisation when the incidence angle is $45--55^o$ and should be better at lower incidence angles, however these results are not followed up in later work.

6.2 Detection strategies

Given that polarimetric SAR is multi-channel data, there are three main detection strategies which can be considered:

A) Apply single channel detectors to each polarimetric channel separately and then fuse the detection results;

B) Fuse the polarimetric channels into a single channel and then apply single channel detectors;

C) Apply multi-channel detection algorithms.

Not much work has been undertaken in the research community to compare these various strategies for ship detection purposes. The only papers where extensive comparisons are made are those by Sciotti *et al.*, [116] and [117]. We discuss the general aspects of that work here, with more specific details in later sections.

For strategy A) Sciotti *et al.*, [116] and [117], fuse the detections by simply combining them together. As single channel detectors they use the cell-averaging CFAR (CA-CFAR) detector (16) which is equivalent to the Power Ratio (PR) detector, the two parameter CFAR (2P-CFAR) detector (15) and the Gaussian Generalised Likelihood Ratio Test (G-GLRT) detector (28). For strategy B), they consider three different methods of fusing

the polarimetric channels into one. The first fusion method is the Minimum Mean Square Error (MMSE) technique, [50], the second is the SPAN technique, [85],

$$\text{SPAN} = |HH|^2 + |HV|^2 + |VH|^2 + |VV|^2, \tag{38}$$

and the third is the Polarimetric Whitening Filter (PWF), [85]. The above single channel detectors are then applied. For strategy C), they only consider the PG-GLRT detector (40) which we discuss in detail below. The performance of these various strategies are compared on simulated SAR data and several L-band SIR-C images. The performance tests are limited and should only be taken as indicative rather than conclusive. In particular, the imagery in [117] is selected to be inhomogeneous in order to emphasise the benefits of segmenting detectors. Good performance is reported for the 2P-CFAR detector when applied to both SPAN and PWF fused imagery. The PG-GLRT detector also performs well.

6.3 Segmenting detectors

Lombardo *et al.* [70] clearly demonstrated the performance gain which segmenting detectors allow on single channel SAR imagery. Consequently, Sciotti *et al.* [117] explored the performance gain which segmenting detectors can achieve on polarimetric SAR imagery. They note that detectors for polarimetric data such as those discussed in the last section suffer from high false alarm rates in inhomogeneous areas. In particular, very dark segments and edges between areas with different mean backscattering cause problems. Segmenting detectors provide a means of overcoming such problems.

Corresponding to the three different detection strategies in section 6.2, Sciotti *et al.* propose three different segmenting detector strategies:

A) Segment each polarimetric channel separately, then apply single channel detectors to each segmented polarimetric channel and then fuse the detection results;

B) Fuse the polarimetric channels, segment the resulting single channel and then apply single channel detectors;

C) Use a multi-channel segmentation algorithm and then apply multi-channel detection algorithms.

The data which Sciotti *et al.* deal with in [117] is single look complex (SLC) quad-pol L-band SIR-C SAR imagery. Therefore they propose using single and multi-channel segmentation algorithms which are based on the "simulated annealing" maximisation of the global likelihood over the image. A reference to the book by Oliver and Quegan, [86] is given. They note that such techniques were especially designed for application to high resolution SAR imagery and therefore are suitable for the SIR-C data. In previous work, [70], such a segmentation algorithm was applied to single channel imagery.

Sciotti *et al.*, [117], report that segmentation of the SPAN-filtered, PWF-filtered and MMSE-filtered data yields equivalent results for ship detection purposes. They also report that in most cases, multi-channel segmentation slightly outperforms (in terms of the segmentation itself) the other methods and rejects small unnatural segments well. However,

over all, they find that the differences between the segmentations produced by the various polarimetric segmentation routines are hardly disceranible by eye. It is of interest to note here that Sciotti *et al.* also considered segmentation based on polarimetric features (sample covariance determinant and entropy factor) but found the resulting segmentations to be less accurate.

As in [70], Sciotti *et al.* adjust the segmentation by performing a raw classification of it into 3 or 4 scattering classes, corresponding to homogeneous areas with different characteristics. The target detectors are then applied within each homogeneous segment with care being taken not to use data from neighbouring segments in the detection feature computation. As single channel detectors, they use the cell-averaging CFAR (CA-CFAR) detector (16), the two parameter CFAR (2P-CFAR) detector (15) and the Gaussian Generalised Likelihood Ratio Test (G-GLRT) detector (28). The only multi-channel detector they use is the PG-GLRT detector (40) which we discuss in detail below. Performance of the various segmenting detection strategies are compared on two different inhomogeneous SIR-C SAR images. They are also compared with the various non-segmenting detection strategies discussed in the previous section. The results show that the segmenting detection strategies allow a substantial improvement over the non-segmenting detection strategies. The best performance is given by applying a 2P-CFAR detector to segmented SPAN-filtered imagery, although for one of the images, the PG-GLRT detector applied to segmented multi-channel imagery also does well. Note however that these performance tests are limited and hence not conclusive.

6.4   The PG-GLRT detector

Motivated by the G-GLRT detector described above in section 5.6, Sciotti *et al.*, [116], developed a similar detector for single look complex polarimetric SAR data. A target window of N pixels is used and the image samples are assumed to be independently and identically distributed from pixel to pixel but correlated within the polarimetric channels. For the n-th pixel, $n = 1, \ldots, N$, the image values over the L polarimetric channels are arranged in to a column vector $\mathbf{x}_n$. Then the vectors $\mathbf{x}_n$ are arranged into a matrix $\mathbf{X} = [\mathbf{x}_1 \ldots \mathbf{x}_n]$. The image samples are assumed to be composed of clutter echoes plus thermal noise under hypothesis (B), and target echoes plus thermal noise under hypothesis (T). The proposed model is that target, clutter and noise echoes are all extracted from zero mean Gaussian random processes with polarimetric correlation properties given by the $L \times L$ matrices $\mathbf{C}_t$, $\mathbf{C}_c$ and $\mathbf{C}_n$. It is further assumed that $\mathbf{C}_t = \text{diag}\{\sigma_{t1}^2, \sigma_{t2}^2, \ldots, \sigma_{tL}^2\}$ and $\mathbf{C}_n = \sigma_n^2 I_L$ where I_L is the $L \times L$ identity matrix. To aid in modelling the background clutter, Sciotti *et al.* also make use of K secondary data $\mathbf{y}_1, \ldots, \mathbf{y}_K$ following hypothesis (B) which are arranged into a matrix $\mathbf{Y} = [\mathbf{y}_{11} \ldots \mathbf{y}_{N1} \ldots \mathbf{y}_{1K} \ldots \mathbf{y}_{NK}]$. With these assumptions, the generalised likelihood ratio in (26) takes the form

$$\frac{\max_{C_t^2, C_c^2, C_n^2} p(\mathbf{X}, \mathbf{Y} | T, C_t^2, C_c^2, C_n^2)}{\max_{C_c^2, C_n^2} p(\mathbf{X}, \mathbf{Y} | B, C_c^2, C_n^2)} = \frac{\max_{C_t^2, C_c^2, C_n^2} \{p(\mathbf{X} | T, C_t^2, C_n^2) \cdot p(\mathbf{Y} | C_c^2, C_n^2)\}}{\max_{C_c^2, C_n^2} \{p(\mathbf{X} | B, C_c^2, C_n^2) \cdot p(\mathbf{Y} | C_c^2, C_n^2)\}} \quad (39)$$

where the unknown parameters are $\sigma_{t1}^2, \ldots, \sigma_{tL}^2$, σ_n^2 and $\mathbf{C}_c$. Substituting the pdfs in and performing the maximisations leads to the detector

$$\frac{|\mathbf{S}_0| \left[\prod_{n=1}^N \left(1 + \mathbf{x}_n^H \mathbf{S}_{n-1}^{-1} \mathbf{x}_n \right) \right]^{K+1}}{\prod_{l=1}^L \left(\sum_{n=1}^N |\mathbf{x}_n(l)|^2 \right)} > \lambda \qquad \Longleftrightarrow \qquad \text{target} \qquad (40)$$

where

$$\mathbf{S}_n = \mathbf{Y}\mathbf{Y}^H + \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i^H$$

This detector is called the Polarimetric Gaussian Likelihood Ratio Test (PG-GLRT) detector. We expect that, just as Lombardo *et al.* in [71] reformulated the GLRT detector into a typical sliding window format, the PG-GLRT detector can also be reformulated into a sliding window format so that the secondary data are taken from a "ring" around the target window.

In [116], Sciotti *et al.* describe the application of the PG-GLRT to both real and synthetic SAR data. They use three polarimetric channels - HH, VV and HV - so $L = 3$. They set $N = 9$ but do not specify the target window arrangement. Presumably a 3x3 pixel sliding target window is used. They also set $K = 3$ but make no comment as to how the secondary data $\mathbf{Y}$ are chosen. However, since the detector theory assumes the samples are independent from pixel to pixel, it is reasonable to take the secondary data from an annulus around the target window as is done in a standard CFAR detector window setup. In fact, it is likely that the PG-GLRT can be developed into a typical CFAR sliding window setup in the same way as Lombardo *et al.*, [71], developed the G-GLRT into a sliding window setup.

The suitability of the PG-GLRT detector for ship detection is assessed by Sciotti *et al.* in both [116] and [117]. In both papers several alternative detection methods are also considered, including applying the cell averaging CFAR detector and the G-GLRT detector to individual polarimetric channels and also to polarimetric data which has been fused into a single channel by methods such as the Polarimetric Whitening Filter. In the first paper, [116], results from experiments with simulated polarimetric SAR data show that PG-GLRT seems to exploit the polarimetric information in an "optimal way" and that it clearly outperforms the other detection techniques. This is backed up with evidence from a simple example with real SAR imagery. The results from the second paper, [117], are not so conclusive. There, performance is tested on two SIR-C SAR images but ones which have been selected for their non-homogeneity. PG-GLRT performs competitively on only one of these images.

6.5   Polarisation entropy

In the article on ocean surveillance by Yeremy *et al.* [142], reference is made to Cloude and Pottier's target decomposition technique for polarimetric SAR data, see [23] and [22]. The technique assumes there is a dominant "average" scattering mechanism in each cell and finds the parameters of this average component. These parameters can then be related to the electromagnetic scattering and physical models of the objects being imaged. For this reason, the method can be used as an unsupervised classification scheme. The technique

operates on the coherency matrix $[T]$ which is derived from the coherent scattering matrix $[S]$ with the use of the Pauli spin matrices. Based on an eigenvalue analysis of $[T]$ two main parameters are extracted — the entropy H and the α-angle, see [23] for details. The entropy H is related to how dominant the dominant scattering component is (a value of zero indicates that all scattering comes from one mechanism and a value of one indicates a completely random scattering mechanism) while the α-angle is related to the type of scattering. Cloude and Pottier, [23], detail a classification scheme in the two-dimensional $H - \alpha$ space. Yeremy *et al.*,[142], claim that preliminary studies have shown the Cloude and Pottier decomposition to be a good method for ship detection. The studies Yeremy *et al.* refer to are in the Satlantic document, [95], by Plache and Henschel. Unfortunately, this report has not been obtained and no further details on the work have been found.

Touzi *et al.* have also considered polarisation entropy for ship detection, see [124], [123], [119] and [118]. Touzi comments that Cloude and Pottier's definition of polarisation entropy is restricted to targets whose degree of polarisation does not change as the transmitted polarisation changes. Hence he introduced a new definition of polarisation entropy, PE, still based on the Von Neumann concept of entropy but which *integrates entropy variations with polarization*. Unfortunately, Touzi does not clearly state what this means or how it is calculated. A thorough search of the literature has shed no light and the author remains uncertain as to exactly how Touzi's polarisation entropy differs from that of Cloude's and how it is calculated. Private communication with Touzi has confirmed that he has not published the details yet but his intention is to do so in the upcoming EUSAR 2004 conference. The only small insight into PE available is a comment in [124] that PE is generally applied with a large window (9x9 pixels). We shall refer to PE as the Touzi polarisation entropy.

As a natural extension of PE, Touzi also introduced Touzi polarisation anisotropy PE^c which is defined by $PE^c = 1 - PE$. Touzi states that polarisation anisotropy characterises target nonstationarity — the higher PE^c is *the larger are signal variations with transmitting-receiving polarizations*. Since the scattering mechanisms of ships are in general more heterogeneous than those of the ocean, ships will have a higher polarisation anisotropy. Thus calculation of polarisation anisotropy is a preprocessing algorithm which enhances ship-sea contrast and it should be followed by some type of bright anomaly detection algorithm. Touzi makes no comment on what sort of detectors are appropriate except for a brief note in [118] that a matched filter was being developed.

In [124], [123], [119] and [118], Touzi *et al.* assess among other things the effectiveness of Touzi's polarisation anisotropy for ship detection. The performance measure used is ship-sea contrast. Unfortunately, this term is not defined in the papers and no indication of how it was measured is given. It is assumed that a quantity akin to the target-clutter-ratio in [142] is meant. The data was collected with the Canadian airborne Convair-580 SAR. However, the motivation behind much of the work was the upcoming launch of polarimetric satellite SARs and hence operational satellite SAR incidence angles ($20--60^o$) were used. The results show that polarisation anisotropy PE^c gives a significant improvement in ship-sea contrast in comparison to single channel polarisations (HH, HV, or VV). Although, it was also found that the effectiveness of PE^c is reduced under rough wind-sea conditions (presumably because the ocean scattering mechanisms become increasingly heterogeneous) and in that case, it is only slightly better than the HV channel alone. It was also found

that at grazing incidence angles (higher than 60^o) PE^c is not effective at all as a means of enhancing ship signatures.

6.6 The Cameron decomposition

The idea behind target decomposition algorithms for polarimetric SAR data is to express the average scattering matrix as a sum of independent elements and to associate different physical scattering mechanisms with the different elements. Further, discussion of target decomposition algorithms may be found in the article [22] by Cloude and Pottier. Also, Yeremy et al. give a good review in [142] of target decomposition for ship detection purposes. Yeremy notes that since target decomposition methods provide information about the scatterer's physical structure, they are suitable for classification as well as detection applications. Here we consider the Cameron decomposition, [13].

Cameron et al. [13] note that two basic physical properties of radar scatterers are reciprocity and symmetry. The scattering matrices of reciprocal scatterers are identified by the property $S_{HV} = S_{VH}$. A symmetric scatterer is defined as a scatterer with an axis of symmetry in the plane orthogonal to the radar line of sight. The scattering matrix S of a symmetrical scatterer is identified by the fact that it is diagonalisable by a rigid rotation transformation, see [52]. The idea behind the Cameron decomposition is to decompose the given scattering matrix into reciprocal and non-reciprocal components and then to further decompose the reciprocal component into a symmetric component which is maximal in some sense and a remaining component. The maximal symmetric component is then assigned to one of six symmetric elemental scattering classes on the basis of the similarity of its normalised form to those of the elemental scatterers. The Cameron decomposition is described by Cameron et al. in [13] and [12] as well as other places. A brief summary is given by Touzi et al. in [120]. We discuss the details presently.

Cloude and Pottier, [22], refer to the Cameron decomposition as a "coherent target decomposition" since it is applied directly to the scattering matrix S. They suggest such decompositions are very susceptible to speckle noise and that this can distort the interpretation of coherent data. They recommend decompositions which are based on the "coherency matrix" instead. However, Yeremy et al., [142], argue that the Cameron method is appropriate for maritime environments. They describe an elemental scatterer as one whose length is small compared to the radar resolution but large compared to the radar frequency. This means the scattering effect from an elemental scatterer is predominantly contained within one resolution cell. They then argue that the backscatter from ship elemental scatterers will be large compared to the surrounding clutter and hence will be the dominant scattering information in ship pixels. They specifically dispute Cloude and Pottier's criticism by arguing that speckle is not a significant factor for ship targets with large backscatter in an open ocean background with low backscatter and little topological variation. They emphasise that highly reflective smooth structures have low surface roughness and hence there is less opportunity for speckle and greater opportunity for coherent backscatter.

The suitability of the Cameron decomposition was also queried by Hawkins et al. in [45] where it is commented that for moderate resolution SAR, each pixel will contain an aggregate of elemental scatterers. However, Ringrose and Harris, [104], and Yeremy et al.,

[142], have had success with the method. Ringrose and Harris actually comment that the technique is not sensitive to the spatial resolution of the input data and is therefore suitable for wide area surveillance imagery. Also, despite their negative comments Hawkins *et al.*, present examples which show that the method has potential and note that with adequate resolution and sufficient study it may be effective. We add here that the work of Hawkins *et al.* appears to be closely related to that of Yeremy *et al.*.

We now describe the Cameron decomposition, [13] and [12], in detail. Our description comes mainly from [13]. The Cameron decomposition was developed with the help of the Pauli matrices:

$$\tilde{S}_a = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \quad \tilde{S}_b = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} \quad \tilde{S}_c = \frac{1}{\sqrt{2}} \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \quad \tilde{S}_d = \frac{1}{\sqrt{2}} \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}.$$

A scattering matrix $\tilde{S}$ is expanded in terms of the Pauli basis by

$$\tilde{S} = \alpha \tilde{S}_a + \beta \tilde{S}_b + \gamma \tilde{S}_c + \delta \tilde{S}_d \tag{41}$$

where and the coefficients α, β, γ and δ may be complex. Clearly the scattering matrices of reciprocal scatterers are exactly those with $\delta = 0$ in this expansion. Thus the first step of the Cameron decomposition is to extract the reciprocal scatterer component of $\tilde{S}$ by setting $\delta = 0$ in (41). Cameron *et al.* denote this component by $\tilde{S}_{rec}$. Cameron *et al.* define an "angle" θ_{rec} which measures the degree to which $\tilde{S}$ satisfies reciprocity. We omit the details here but we note that in the original paper, [12], scattering matrices for which $\theta_{rec} > \pi/4$ are considered to be from non-reciprocal scatterers and are excluded from further analysis.

Next, Cameron *et al.* seek to extract the maximal symmetrical scatterer component of $\tilde{S}_{rec}$. In order to best define what this means, the problem is vectorised. This is done with the operator

$$\mathbf{V}: \quad \tilde{S} = \begin{bmatrix} s_{11} & s_{12} \\ s_{21} & s_{22} \end{bmatrix} \quad \mapsto \quad \vec{S} = \begin{bmatrix} s_{11} & s_{12} & s_{21} & s_{22} \end{bmatrix}^T. \tag{42}$$

Cameron *et al.* denote the norm of a vector $\vec{u}$ by $||\vec{u}||$, the inner product of two vectors $\vec{u}$ and $\vec{v}$ by $(\vec{u}, \vec{v})$ and uses the notation $\hat{u}$ to indicate that $\vec{u}$ is a unit vector. Thus the Pauli basis becomes $\hat{S}_a = \mathbf{V}\tilde{S}_a$, $\hat{S}_b = \mathbf{V}\tilde{S}_b$, $\hat{S}_c = \mathbf{V}\tilde{S}_c$ and $\hat{S}_d = \mathbf{V}\tilde{S}_d$.

Unfortunately, while the set of all scattering matrices of symmetrical scatterers is subset of the subspace of reciprocal scatterer scattering matrices, it is not a subspace itself which complicates the analysis. However, Cameron *et al.* show that the vector form of any such matrix $\tilde{S}_{sym}$ has an expansion of the form

$$\vec{S}_{sym} = \alpha \hat{S}_a + \epsilon[\cos(\theta)\hat{S}_b + \sin(\theta)\hat{S}_c] \tag{43}$$

where $\theta = -2\psi$ and ψ is the rotation angle which diagonalises $\tilde{S}_{sym}$ (recall here that the scattering matrix of a symmetrical scatterer is identified by the fact that it is diagonalisable by a rigid rotation transformation, [52]). Thus, in seeking the symmetric component of $\vec{S}_{rec}$, Cameron *et al.*, [12], consider decompositions of the form

$$\vec{S}_{rec} = \vec{S}_{sym} + \vec{S}_2 \tag{44}$$

for some $\vec{S}_2$ where

$$\vec{S}_{sym} = (\vec{S}_{rec}, \hat{S}_a)\hat{S}_a + (\vec{S}_{rec}, \hat{S}')\hat{S}' \tag{45}$$

and

$$\hat{S}' = \cos(\theta)\hat{S}_b + \sin(\theta)\hat{S}_c \tag{46}$$

and $(\cdot, \cdot)$ denotes the inner product. To guarantee that $\vec{S}_{sym}$ is the largest symmetric component which can be extracted from $\vec{S}_{rec}$, they choose the value of θ in (46) so that it maximises $|(\vec{S}_{rec}, \hat{S}')|$. We denote the maximal symmetric component derived in this way by $\vec{S}_{sym}^{max}$. While our interest lies entirely in $\vec{S}_{sym}^{max}$, it is of interest to note here that according to Cameron *et al.*, the remaining component is also a symmetrical scatterer scattering matrix which is orthogonal to $\vec{S}_{sym}^{max}$.

It is shown in the appendix of [12] that the value of θ in (46) which generates $\vec{S}_{sym}^{max}$ satisfies

$$\tan(2\theta) = \frac{\beta\gamma^* + \beta^*\gamma}{|\beta|^2 + |\gamma|^2} \tag{47}$$

unless $\beta\gamma^* + \beta^*\gamma = 0 = |\beta|^2 + |\gamma|^2$ in which case $\theta = 0$. By computing the terms in (45) we also find that $\vec{S}_{sym}^{max}$ is of the form given in (43) with

$$\epsilon = \beta\cos\theta + \gamma\sin\theta. \tag{48}$$

Here α, β and γ are the coefficients of the original Pauli expansion (41). As noted above, (45) is diagonalised by a rotation through an angle $\psi = -\theta/2$ hence (as noted by Touzi *et al.* in [120]) after diagonalisation

$$\vec{S}_{sym}^{max} = \alpha\hat{S}_a + \epsilon\hat{S}_b \tag{49}$$

The classification scheme of the Cameron decomposition, [12], makes use of the angle τ defined by

$$\cos\tau = \frac{|(\vec{S}_{rec}, \vec{S}_{sym}^{max})|}{||\vec{S}_{rec}||\,||\vec{S}_{sym}^{max}||}. \tag{50}$$

which measures the degree to which $\vec{S}_{sym}^{max}$ deviates from $\vec{S}_{rec}$. Scattering matrices for which $\tau > \pi/8$ are considered to be from asymmetrical scatterers and are classified separately.

The final step of the Cameron decomposition involves specifying a normalised diagonal form for the scattering matrix of a symmetrical scatterer and thereby associating a physical scattering mechanism with the matrix $\vec{S}_{sym}^{max}$. The normalised form is obtained by scaling and rotating the scattering matrix and is given by

$$\hat{\Lambda}(z) = \frac{1}{\sqrt{1+|z|^2}} \begin{bmatrix} 1 \\ 0 \\ 0 \\ z \end{bmatrix}, \qquad z \in \mathbf{C}, \qquad |z| \le 1. \tag{51}$$

Thus the parameter which determines the scatterer type is z. Of course the rotation angle which diagonalises $\vec{S}_{sym}^{max}$ is $\psi = -\theta/2$ where θ satisfies (47). However, care needs to be taken here. Cameron *et al.* note that rotations through $\psi \pm \pi/2$ will also diagonalise $\vec{S}_{sym}^{max}$ although with the diagonal elements interchanged. Depending on the scattering matrix,

Scatter type	Normalised form
Trihedral	$\hat{\Lambda}(1)$
Dihedral	$\hat{\Lambda}(-1)$
Dipole	$\hat{\Lambda}(0)$
Cylinder	$\hat{\Lambda}(1/2)$
Narrow Diplane	$\hat{\Lambda}(-1/2)$
Quarter Wave Device	$\hat{\Lambda}(i)$

Table 4: Cameron's elemental scatterer types

there may be further ambiguities in the rotation angle. Cameron *et al.* resolve these ambiguities to ensure $|z| < 1$ in (51). It follows that the scatterer type associated with a particular scattering matrix can be represented by a point z in the unit disc of the complex plane.

As reference models of scatterer types, Cameron *et al.* consider the symmetric scattering mechanisms shown in Table 4. In order to compare general scattering matrices of symmetric scatterers with these reference models, Cameron *et al.*, [13], define the metric

$$d(z_1, z_2) = \cos^{-1}\left(\frac{\max(|1 + z_1 z_2^*|, |z_1 + z_2^*|)}{\sqrt{(1 + |z_1|^2)}\sqrt{(1 + |z_2|^2)}}\right). \tag{52}$$

where $\hat{\Lambda}(z_1)$ and $\hat{\Lambda}(z_2)$ are the normalised diagonal forms of the scattering matrices. The scatterer type of a general symmetric scattering matrix is then assigned to the scattering mechanism in Table 4 which is closest according to this metric. We note here that in earlier work, [12], Cameron *et al.* developed a complete classification scheme for scattering matrices based on the above decomposition but the method for comparing symmetric scatterer scattering matrices was different and did not involve the metric (52). We also note that Touzi *et al.*, [120], suggest an improvement to the metric (52). Touzi *et al.* point out that Cameron *et al.* derived the metric (52) prior to diagonalising both scattering matrices and that it makes more sense to to take the magnitude of the complex scalar product of the normalised scattering matrices $\hat{\Lambda}(z_1)$ and $\hat{\Lambda}(z_2)$ as the metric. Thus they propose the metric

$$d(z_1, z_2) = \frac{|1 + z_1 z_2^*|}{\sqrt{(1 + |z_1|^2)}\sqrt{(1 + |z_2|^2)}} \tag{53}$$

which is simpler than (52).

The first use of Cameron's decomposition for ship detection purposes was by Ringrose and Harris, [104]. This work is really only a feasibility study but it obviously has the potential to be developed into a valid ship detection algorithm. Ringrose and Harris applied the method to SIR-C SAR imagery of the ocean and simulated SAR imagery of ships. The SIR-C data was L-band single look complex fully polarimetric imagery. The simulated ship data used the same radar parameters as the SIR-C SAR imagery, the ship was 50 m in length and resolutions of 25 m and 100 m were used. They found that for the most part, the scatterer types in ship pixels differ from those in ocean pixels and hence ship pixels can be easily detected. However, they also used a threshold in order to eliminate dull pixels from consideration (ship pixels are usually bright). They note that

the threshold is not really linked to a false alarm rate and so there is no trade off between false alarm rate and detection probability in their method.

As an example of the results produced by Ringrose and Harris in [104], with a threshold of -10 dBm the only ocean scatter types present are cylinders, trihedrals and dipoles. Three different ocean areas are analysed and other thresholds are considered. The main ocean scatterer type is a cylinder. Interestingly, Ringrose and Harris find that wind speed does not seem to affect the ocean scatterer types. The variation of ship scatterer type with aspect angle is also investigated. From this, Ringrose and Harris conclude that the most important scatterer type for ship detection is the dihedral (or narrow dihedral). However, they find that at certain aspect angles the ship scatterers do not differ from the ocean scatterers and so their approach will not lead to a perfect ship detection method. On the other hand, they comment that the method is advantageous since resolution is not critical and hence should work with low resolution wide area surveillance imagery.

The Cameron decomposition was also investigated for ship detection purposes by Yeremy *et al.* in [142]. Like the work of Ringrose and Harris, the Yeremy study is explorative rather than developmental. However, the Yeremy study involves a much larger data set and is much more comprehensive. It warrants a careful reading by anyone wishing to develop a ship detection algorithm based on the Cameron decomposition. We provide a brief overview of the relevant parts of the Yeremy paper in the following paragraphs.

The data Yeremy *et al.* use has been described above in the section on optimal polarisations for ship detection. It was collected by CCRS Convair-580 C-band SAR system during the MARCOT'98 and CRUSADE'00 trials. More than 75 ships were selected from the data by manual visual means. The ships were extracted using visual delineation of the target area or supervised thresholding. Each ship was centred in a 800x100 image chip with land, other targets and saturated pixels masked out. The resulting "imagettes" were then used to assess detection performance of the Cameron method. A valid detection was recorded for the whole imagette if at least one ship pixel was detected and no ocean pixels were detected.

Like Ringrose and Harris, Yeremy *et al.* extend the basic Cameron method by including a threshold to take account of the backscattered magnitudes. There are at least two arguments in the Yeremy paper justifying this: first, the ocean may have similar scattering mechanisms to ships and the backscattered intensity will help differentiate the two; and second, since the Cameron method is susceptible to speckle, only large backscatter pixels are desired anyway. Consequently, in the method Yeremy *et al.* implement, pixels are only considered detected if they are above the chosen threshold. Yeremy *et al.* note that ideally the choice of threshold should be based on a model (empirical or theoretical), but they know of no validated polarimetric threshold models and suggest that further research using ground-truthed data is needed on this topic. Instead for demonstration purposes only, the chosen threshold was determined from the data empirically by undertaking a target-to-clutter ratio study. On this basis it was decided to set the threshold to be 10 dB less than the maximum SPAN (see (38)) of the imagette being processed. Yeremy *et al.* state that the combination of threshold and Cameron decomposition gives better detection rates than either method individually.

The original Cameron classification scheme, [12], could identify eleven different elemental scatterers. Of these, five were chosen as likely to be indicative of man-made vessels

and the study was limited to them. The selected scatterer types were: dihedral, narrow diplane, cylinder, diople and quarter wave device. The results for each of these scatterer types were presented separately in order to emphasise which scatterers are significant for ship detection purposes. Trihedral scatterers were not chosen since it was believed that trihedral shapes are not common on ships. Yeremy *et al.* note here that the Ringrose and Harris study found that ships are dominated by dihedral, narrow diplane and quarter wave device scatterers.

Yeremy *et al.* note that when using the metric (52) to compare the parameter z_1 of an image pixel scattering matrix, with the parameter z_2 of an elemental scatter scattering matrix, if $d(z_1, z_2) = 0$ then the data and the model are the same whereas if $d(z_1, z_2) = \pi/2$ then the data and the model are orthogonal and hence very dissimilar. For this reason, Yeremy *et al.* add the constraint that $d(z_1, z_2) < \pi/12$ since this ensures that only dominant elemental scatterers are used in the analysis.

Detection results for the various Cameron elemental scatterers were good, and in particular, the narrow diplane and quater wave device elemental scatterers gave detection rates of 98% and 91% respectively. Further, by combining the results for several scatterer types even greater confidence levels could be obtained with 100% detection rate being recorded for several combinations. Yeremy *et al.* conclude that although the Cameron method has been criticised because it relies on single pixel operations, their results indicate that it is a very robust ship detection method (as long as only dominant scatterers are used). Further discussion of these results may be found in [143].

Two other points in the paper are worth highlighting here. First, Yeremy *et al.* observe that the spatial density of scatterer types may be an important factor in ship detection — for instance the spatial density of detected narrow diplane scatterers increases at the ship location relative to the surrounding clutter whereas the reverse is true of detected cylindrical scatterers. Second, six of the imagettes used in the study produced anomalous results for the Cameron method. Yeremy *et al.* believe that these imagettes contained false targets and that the targets were in fact ocean buoys. On this basis, their results indicate the classification information in polarimetric methods such as the Cameron decomposition may be able to discriminate false targets from true targets.

In the final part of the paper, Yeremy *et al.* explore the classification potential of the Cameron method. The idea here is to use the spatial composite of the detected elemental scatterers as classification features. Yeremy *et al.* found that each ship in the study had individual characteristics in terms of detection types, spatial densities and distributions. However, they add that further research and more ground-truthed data are required. Several example ship signatures are presented. For one of these, a photograph of the original ship was also available and interesting and convincing comparison is made between the detected elemental scatterers and the various physical structures on the ship.

6.7   Symmetric scattering characterisation method (SSCM)

In follow-up work, Touzi and Charbonneau have suggested several improvements to the Cameron decomposition, [120] and [121]. They also extract several new parameters from the decomposition. While their application is to ship characterisation and identification,

[124] and [122], the method remains essentially the same and hence should be relevant to ship detection. For this reason it is appropriate to review their work here. We emphasise however that these papers are investigative and suggest avenues for further exploration rather than document polished algorithms. We also add that this makes the papers difficult to follow and the author's meaning is not always clear.

In [124] and [120], Touzi *et al.* distinguish two types of target decomposition methods: coherent target decomposition (CTD) methods and partially coherent target decomposition (PCTD) methods. They note that PCTD methods (such as the Cloude and Pottier method, [22] use the multi-look Mueller, covariance or coherency matrix and therefore are low resolution methods. On the other hand CTD methods (such as the Cameron decomposition) extract polarisation information directly from the 1-look scattering matrix and are therefore high-resolution methods. High-resolution methods are desirable for ship classification and hence the Cameron decomposition is of interest.

Touzi *et al.*, [124] and [120], have two main criticisms of the Cameron method. First, they state that Cameron's method uses a very coarse classification scheme. They are referring here to the use of the metric (52) to classify a normalised scattering matrix, $\hat{\Lambda}(z)$, as one of the elemental scatterer types $\hat{\Lambda}(z_{\text{ref}})$. They note that the ratio $|z|/|z_{\text{ref}}|$ (expressed in decibels) can be unacceptably large and that Cameron's classification can therefore be misleading. By analogy with SAR channel imbalance and known polarimetric SAR radiometric calibration requirements, they suggests that a threshold of 1 dB be applied to the ratio in order to limit "radiometric dispersion". However, with the threshold at this level they find that too many pixels remain unclassified and therefore Cameron's classification is of limited use. Their resolution to this dilemma is to avoid the classification scheme entirely and in effect they work directly with the parameters of the normalised scattering matrix. As an example, they note that the Cameron method leads one to believe that ocean clutter is dominated by cylindrical scattering and ship clutter by dipole and narrow dihedral scatterers, [104] and [142]. However, they show that when the coarse classification scheme of the Cameron method is dropped it becomes apparent that less than 1% of ocean pixels are pure cylindrical scatterers.

Their other criticism of the Cameron method is that it is a CTD method yet it does not include an assessment of scattering coherence and can therefore be applied in areas of non-coherent scattering, leading to meaningless results. They state that *CTD can only be applied under coherency conditions in which the polarimetric information is fully represented by the scattering matrix* and further, *if the target has a coherence time that is significant in comparison to the measurement time, the scattered wave is then partially polarised and cannot be represented with the scattering matrix.* Consequently, they recommend that the SAR image being processed be segmented into coherent and non-coherent scattering classes and that CTD methods be applied within areas of coherent scattering and PCTD methods be applied within areas of non-coherent scattering. We note here that the ocean surface is continually changing shape and so is likely to genereate non-coherent scattering whereas ships are more likely to generate coherent scattering.

They propose two ways of assessing the scattering coherence: one for point targets and one for distributed targets. The point target coherent scattering test is designed to locate scattering which is coherent within a single resolution cell. Scattering is considered to be coherent if the resolution cell is dominated by an elemental coherent scatterer. Thus

a Rician scattering model (fully developed speckle dominated by a constant phasor) is used. The Rician model is parameterised by its signal-to-clutter ratio. Resolution cells for which this ratio is above a 15 dB threshold are considered to contain coherent scattering. Details are given as to how the signal-to-clutter ratio can be deduced either from either a Huynen decomposition or a Graves matrix decomposition. The distributed target coherent scattering test is designed to locate scattering which is coherent over many resolution cells but possibly only locally time-wise coherent in each. This is done through the use of a quantity p_{sym} (see equation (56)) which measures symmetric scattering coherence. A coherence map is generated with a moving window and pixels with p_{sym} close to 1 are taken to represent areas of (locally) coherent scattering. Touzi *et al.* do not explicitly specify a window size or a threshold value for p_{sym}.

A full description of the SSCM is given by Touzi and Charbonneau in [120]. We go through the details next. The main steps are:

- Calculate the parameters α and ϵ of the maximum symmetric component according to (41), (47) and (48);

- Calculate d_{sym} according to (54) or (50) and use it to assess the significance of the symmetric scattering component;

- (Distributed target coherence test) Segment the image into coherent and non-coherent classes using the map of p_{sym} as calculated by (56);

- (Point target coherence test) Apply the radiometric Rician threshold within the noncoherent class — pixels with a signal-to-clutter ratio above the 15 dB threshold are assigned to the coherent class

- Generation and analysis of the $\vec{S}_{sym}^{max}$ parameters within the coherent class

We explain d_{sym}, p_{sym} and the $\vec{S}_{sym}^{max}$ parameters next.

Touzi and Charbonneau use $\cos \tau$ in (50) as a measure of the size of symmetrical scattering component of the scattering matrix $\vec{S}$. Hence they define $d_{sym} = \cos \tau$ refers to it as the degree of symmetry. They note that d_{sym} is in fact the ratio of the intensity of $\vec{S}_{sym}^{max}$ to that of $\vec{S}_{rec}$ and hence

$$d_{sym} = \frac{|\alpha|^2 + |\epsilon|^2}{|\alpha|^2 + |\beta|^2 + |\gamma|^2}. \tag{54}$$

In the application of SSCM in [120], Touzi and Charbonneau calculate a map of d_{sym} using a 5x5 pixel sliding window. Thus the parameters in (54) are replaced by the appropriate expected values over the window. They exclude asymmetrical scatterers from the analysis by applying a lower threshold of 0.92 to d_{sym}. In other words, pixels with $d_{sym} < 0.92$ (corresponding to $\tau > \pi/8$) are excluded.

In [13], Cameron *et al.* commented that the unit disc representation of $\vec{S}_{sym}^{max}$ given by the normalised form in (51) may not be the best representation. As an alternative,

Touzi and Charbonneau suggest that the Poincare sphere be used. Hence they propose the normalised form

$$\hat{\Lambda}(\psi_c, \chi_c) = \begin{bmatrix} 1 \\ \cos(2\psi_c)\cos(2\chi_c) \\ \sin(2\psi_c)\cos(2\chi_c) \\ \sin(2\chi_c) \end{bmatrix}. \tag{55}$$

where ψ_c and χ_c are coordinates of the Poincare sphere. Touzi *et al.* outline how ψ_c and χ_c can be calculated from the parameters α and ϵ of $\vec{S}_{sym}^{max}$ in (49) but the details are beyond the scope of this report. There is a rotation phase ambiguity here which they resolve by restricting ψ_c to the interval $[0, \pi/2]$ implying that only half the sphere is used. While they determine where the elemental scatterer types in table 4 lie on the sphere, they do not define a metric like that in (52) and nor do they propose an associated classification method. Instead, they work with the parameters ψ_c and χ_c as they stand.

Touzi and Charbonneau note that only coherent symmetric scatterers can be represented by points on the surface of the sphere and that partially coherent symmetric scatterers are represented by points inside the sphere. They state that *a partially coherent symmetric scatterer is represented as a point inside the sphere at a distance from the sphere centre determined by the degree of coherence of the scatterer components α and ϵ on the basis $(\hat{S}_a, \hat{S}_b$ given by*

$$p_{sym} = \frac{\sqrt{(\langle|\alpha|^2 - |\epsilon|^2\rangle)^2 + 4|\langle\alpha \cdot \epsilon^*\rangle|^2}}{\langle|\alpha|^2 + |\epsilon|^2\rangle}. \tag{56}$$

When symmetric scattering is dominated by the trihedral, $\epsilon = 0$, or the dihedral, $\alpha = 0$, the degree of coherence p_{sym} of the symmetric scattering is close to 1 as expected. This expectation motivates the choice of p_{sym} to measure the degree of coherence... Note the use of expected values in (56). In the application of SSCM, Touzi and Charbonneau calculate a map of p_{sym} using a 5x5 pixel sliding window. In several places they use a threshold of 0.8 on p_{sym} but this value is not explicitly recommended anywhere.

We turn now to the parameters which Touzi and Charbonneau extract from $\vec{S}_{sym}^{max}$. We have already discussed the Poincare sphere coordinates ψ_c and χ_c. Touzi *et al.* also extract an angle η and a phase difference $\phi_a - \phi_b$. This is done by re-expressing $\vec{S}_{sym}^{max}$ from the form (49) to the form

$$\vec{S}_{sym}^{max} = \exp j\phi_a \cdot \sqrt{|\alpha|^2 + |\epsilon|^2} \cdot [\cos(\eta) \cdot \hat{S}_a + \sin(\eta)\exp j(\phi_a - \phi_b) \cdot \hat{S}_b] \tag{57}$$

where η lies in the interval $[0, \pi/2]$. Touzi and Charbonneau note that η is analogous to Cloude and Pottier's parameter α, [22] and [23]. They comment later in the paper that partitioning η into the three classes $\eta = 0$, $\eta = \pi/2$ and $0 < \eta < \pi/2$ is equivalent to Van Zyl's odd-even bounce coherent scattering classification. They suggest that both η and $\phi_a - \phi_b$ will contain useful information. They also mention that the channel coherence $\gamma_{\hat{S}_a, \hat{S}_b}$ should be of interest.

In the final section of the paper, [120], Touzi and Charbonneau discuss the application of the above methods to an 18 m ship imaged by the Canadian Convair-580 C-band SAR. Their results clearly demonstrate the effectiveness of the SSCM. The parameters, ψ_c, χ_c, η and $\phi_a - \phi_b$ are all shown to carry useful information. In the conclusion of the paper,

they make the following comment. *The preliminary results obtained with the Convair-580 dataset show that the symmetric component of target scattering has a big potential for ocean characterisation and ship identification.* In a follow up paper, [124], Touzi *et al.* apply the SSCM to SAR imagery of two ships produced by the Convair-580 system. The results validate the potential of SSCM for ship characterisation and demonstrate in particular that ship orientation (pitch and roll) can be estimated using SSCM. This was done by applying the SSCM to the ship mast and other structures which exhibited symmetric scattering. It was noted that on the one hand, the SSCM was reasonably robust to rough sea-wind conditions but on the other, SAR focus errors led to errors in estimating the symmetrical scatterer properties. The conclusion was that SSCM can be used to provide a ship specific distribution of "permanent" scattering targets and hence it may be useful for classification and identification purposes

7   Prescreening — interferometric SAR

Along-track interferometric SAR (InSAR) allows the enhanced detection of moving targets on the ocean surface including ships and ship wakes. In along-track InSAR systems, two separate radar antennae are mounted on the same aircraft but displaced along the body of the aircraft in the direction of flight. The signal received by each antenna is processed to form a single look complex SAR image and then an interferogram is formed from them by multiplying one image by the complex conjugate of the other. The magnitude of each pixel in the interferogram is related to the RCS of the corresponding resolution cell while the phase is related to its average velocity component in the radial direction. With a properly calibrated system, the relationship between interferogram phase and radial velocities can be quantified and accurate velocity measurements made.

Campbell *et al.* describe the Canadian Centre for Remote Sensing (CCRS) along-track InSAR system in [15]. They quantify the relationship between interferogram phase and radial velocities as follows. Suppose the antennae are separated by a distance d and the aircraft has a velocity V. Then the time delay between the formation of the two SAR images is $t = d/V$. In that time, a target with radial velocity v will have moved a distance vt in the radial direction and hence its phase will have changed by $4\pi vt/\lambda$ where λ is the radar wavelength. If Φ is the interferogram phase it follows that

$$v = V\lambda\Phi/2d\pi \qquad (58)$$

where the reduction by a factor of two is due to the fact that the same transmit antenna is used for both InSAR antennae. Campbell *et al.* note that for the CCRS system t is typically around 5 ms whereas the ocean decorrelation time is thought to be between 50 and 100 ms.

Campbell *et al.*, [15], note that the velocity measurements provided by along-track In-SAR are independent of the RCS. Consequently, InSAR allows the detection of small ships which may not be much brighter than the background ocean clutter. If such ships have a velocity component in the radial direction which is different to that of the surrounding ocean, they will show up as anomalies in the velocity image. An added advantage is that one can estimate true velocities from radial velocities if the ship's heading (which can be

determined from either its shape or its wake) is known. Campbell *et al.*, [15], demonstrated the ability of along-track InSAR to detect moving ships and ship wakes during the 1995 MARCOT trial off the coast of Novia Scotia. However, no automatic processing of the imagery was undertaken.

A similar approach is used by Schulz *et al.* in [113] and [112] although their interest is in the special problem of detecting slow moving cargo ships in locks. Schultz *et al.* include in the formula (58) a term to account for the radar incidence angle and hence produce ground velocities rather than radar line of sight velocities. The main thrust of their papers is to deal with problems caused by areas of low signal-to-noise ratio in the phase information. Their solution is to use intensity information to identify and remove such areas. They also use a region growing method to cluster pixels with similar velocity estimates.

A somewhat different technique of using SAR interferometry for ship detection has been described by Arnaud in [6]. Again, the idea is to use phase rather than intensity to aid in the detection of small ships with low RCS. However, Arnaud proposes formation of the interferogram from two different looks processed from a single aperture SAR system. This is done by selecting two different non-overlapping bands from the synthetic aperture Doppler spectrum. Of course, the resulting images have reduced resolution but they also correspond to imaging the scene from two slightly different angles at different times. It is the time difference that is critical to Arnaud's application. His application is to the ERS satellite SAR in which case the time difference will be hundreds of milli-seconds. Arnuad's presumption is that during this time delay the sea which is continuously moving will have strongly incoherent behaviour while ships, like point scatterers will have coherent behaviour.

Arnaud's scheme is therefore to: extract two non-overlapping Doppler bands from the raw SAR data; form the SLC images for each look; and then form an interferogram and its "phase coherence image". Pixels with high "phase correlation" are expected to be ship pixels. Arnaud detects them by applying a simple threshold to the phase coherence image. Arnaud has applied this method to ERS-1 SAR data with promising results. However, the tests were not extensive and the imagery was not ground-truthed so although feasible, the practical utility of the method has not been demonstrated. Arnaud notes that the method can be fine tuned and some improvements may result from considering the following: the width of the Doppler subbands used; the spectral separation of the Doppler subbands used; the method of computation of the coherence image (and especially the window size used). On the first two points, Arnaud comments that both the separation and subband widths need to be increased in calm seas to enhance the decorrelation of the sea whereas in agitated seas they need to be reduced to avoid the decorrelation of the ship.

Ideas very similar to Arnaud's have been developed in the work of Iehara *et al.*, [53], and Ouchi *et al.*, [93], — the second paper being an extension of the work in the first. They likewise suggest taking advantage of the fact that two different looks processed from different sub-apertures of a SAR system will be separated by a small time delay. They argue that owing to this time delay and the different scattering mechanisms, there will be little correlation of the ocean backscatter in the two images but there will be a large correlation in the backscatter from ships. Thus they propose using a measure of the correlation between the looks to detect ships.

The basic processing flow, as outlined in the first paper by Iehara *et al.* [53], is: generate the split-look images from the raw SAR data; extract corresponding windows from each split-look image; compute the two-dimensional cross correlation function (2D-CCF) for the windows; if the 2D-CCF has a peak, the windows are judged to contain a ship; process the entire image by moving the windows around. The authors note that the method is not affected by ship motion because the only effect of ship movement is to shift the location of the peak in the 2D-CCF. The method was tested on Radarsat imagery but not extensively.

In the follow up paper [93] by Ouchi *et al.*, the method is further tested. One of the advantages of the method is that it can detect ships even when the intensity of the ship backscatter is similar to that of the surrounding ocean. Consequently, Ouchi *et al.* are particularly interested in that case. While the first few steps of the algorithm are the same as in [53], it appears that a different technique is used to calculate the correlation between the two windows, with a coherence value being produced rather than a cross-correlation function. Unfortunately, the paper is poorly written and the details are not clear. Be that as it may, the final result of the correlation process described by Ouchi *et al.* is a "coherence image" which is then thresholded to detect ships. Ouchi *et al.* use a simulation study to identify the optimum size of the window used to measure the correlation between the looks and then they evaluate the method on Radarsat data. They find that an 8 x 8 pixel window is best but that ships smaller than 8 x 8 may not be detected. Their experiments with modified Radarsat data verify that the method is feasible and can detect otherwise invisible ships. It is noted that other ocean features such as current boundaries and breaking waves can generate false alarms and that further processing may be needed to eliminate them.

8   Discrimination algorithms

We begin this section with a general introduction to the philosophy behind ship target discrimination. We follow that with a general discussion as to how ship wake detection can be used to aid ship target discrimination. We then observe that not much research has been undertaken on discrimination algorithms for ship detection and offer reasons as to why. Following that we discuss the common but elementary discrimination algorithms based on simple target measurements. We then review the Ocean Monitoring Workstation's approach of using human supervision for target discrimination. The remainder of the section consists of a detailed review of the various discrimination algorithms documented in the literature.

Most ship detection systems have a discrimination phase following the prescreening stage. The idea of this stage is to reduce the false alarm rate. The algorithms in the discrimination stage process the candidate targets detected by the prescreening phase and are aimed at discriminating true targets from false alarms. To do this, discrimination algorithms usually operate on an image chip which contains the candidate target pixels and their local surroundings. They may or may not take account of which pixels were detected by the prescreeners. The algorithms which do not take account of prescreening detections have the potential to be prescreening algorithms themselves. Conversely, prescreening

algorithms also have the potential to be discrimination algorithms. The distinction becomes one of speed and accuracy. It is generally the case that prescreening algorithms are fast and therefore less accurate while discrimination algorithms are more accurate and therefore slow. We add here that the distinction between prescreening and discrimination algorithms can often be blurred with discrimination algorithms incorporated into the prescreeners for efficiency and other reasons. Also, just as preprocessing and prescreening algorithms are linked, it is to be expected that there will be some dependency between prescreening and discrimination algorithms. A good discriminator for one prescreener may not work well with another.

As an example of the linking between prescreening algorithms and discrimination algorithms we cite the work of Ferrara and Torre, [34]. They consider two different false alarm rates (FAR) in the design of their prescreening algorithm. The lower FAR is for the case where no restriction is placed on the detected pixels while the higher FAR is for the case where isolated detected pixels are considered too small to be ship pixels and are discarded. This latter case is a very simple type of discrimination algorithm. (Of course its validity depends on the surveillance scenario at hand and on image resolution and ship size in particular.) The point is that by improving false alarm control, the prescreener FAR can be increased and hence better probabilities of detection can be obtained.

Many ship detection systems include a wake detection and analysis module following ship detection. However, while wakes are an obvious additional source of information about ships, their use in target discrimination algorithms is limited. The problem is that stationary and slow moving ship have no wakes. Further, wakes are unreliable or not visible in many types of SAR imagery. Thus although the presence of a wake confirms the presence of a ship, the converse is not true and the absence of a wake does not imply the absence of a ship. Consequently wakes can really only be used to increase the confidence in detections rather than reject false alarms. However, having said that, we note that a reasonable compromise is to retain all large detections but reject small (or otherwise questionable) detections unless a wake is present. This is the approach taken by Eldhuset in [32]. A variation of this is proposed by Jiang *et al.* in [55]. The authors suggest two prescreener thresholds be used — detections above the higher threshold are retained as usual (subject to the chosen discrimination algorithms) while detections above the lower threshold but below the higher threshold are only retained as long as a corroborating wake can be found.

In addition to confirming detections, wakes can also provide ship course and speed information. Such information is ancillary to the detection and is not usually for discrimination/classification purposes. The ship course is determined by the orientation of wake and the ship speed can be deduced from the azimuth offset between the wake and the ship using formulae (5) and (6). Of course, the later formulae only provide the azimuth component of the speed and the ship's orientation needs to be taken into account in order to estimate the true velocity. Also, Eldhuset, [32], warns that speed estimation is unreliable if the angle between the azimuth direction and the ship orientation is larger than 15^o (else the range component of the ship's velocity is too small). We also note that wakes can be both bright and dark and can have several components at varying angles making analysis and orientation estimation difficult. As mentioned in the introduction, wake detection is beyond the scope of this report and we will only provide a cursory discussion here. We will indicate which systems include wake detection and how it is used without going into

the details.

With the exception of wake detection and analysis there has been little research effort on target discrimination algorithms for ship detection. The most likely explanation for this state of affairs is the lack of well ground-truthed data. We have already noted the difficulty in producing good theoretical models of target backscatter given the large number of variables which come into play. Consequently, discrimination algorithms often use empirically derived target and background models. Deriving such models requires a large amount of training data. By training data we mean many well ground-truthed examples of true ships and false alarms to learn from. Moreover the examples need to be representative of the context in which the discrimination algorithm will be applied. Even in the case of discrimination algorithms which do not need training data, well ground-truthed test data is still needed for parameter setting and performance assessment. Such databases are distinctly lacking. See for instance: the comment by Rey *et al.* on this in the conclusion of their report, [103]: the comment by Yeremy *et al.* in [142] that rarely have there been ship detection studies with more than three ground-truthed ships; and the comment by Yeremy *et al.* in [143] that the CRUSADE'00 trial is an invaluable set of data since few other experiments have been as extensive.

Another likely explanation for the lack of research into discrimination algorithms is that most of the applications are to wide area surveillance. Consequently, low resolution, wide swath data is used. The prime example is Radarsat ScanSAR imagery with a resolution of 200 metres and a pixel spacing of 100 metres. Given that few ships are more than a few hundred metres long and most are shorter, only a handfull of image pixels will contain the ship backscatter and many of them will also contain ocean backscatter. For smaller ships, it is possible that the backscatter is contained in a single pixel. It is hard to extract worthwhile discrimination features from such scant data.

For single channel data, commonly all that is done for target discrimination is to cluster contiguous detected pixels into a single detection and then extract some simple physical measurements from the target. The target measurements include: area, length, width, total intensity (i.e. RCS), peak intensity, orientation, centroid coordinates, aspect ratio, etc. Length, width and orientation can be estimated by total least squared error regression, principal component analysis or fitting a circumscribing rectangle. Note that estimated ship dimensions are not always accurate owing to sidelobe effects, smearing caused by ship motion, spectral leakage, etc. Bounds are placed on the target measurements to eliminate detections which are likely to be false alarms or otherwise not of interest. For instance it is common to remove detections which are too small and therefore likely to be due to the effects of speckle noise or local wind effects. It is also common to remove detections which are too large. For instance, this is done by Kourti *et al.*, [61], when searching for fishing vessels — detections with more than 5 pixels in a row or column are separated out since the SAR resolution is 50 to 70 metres and fishing vessels are shorter than 150 m and therefore are unlikely to generate such detections. Target measurements can also be used for classification purposes. An example of this is given by Askari and Zerr in [7] where it is noted that the length-to-width ratio is a standard indicator of ship structure. For merchant ships this ratio is of the order of 6 to 7, whereas it is larger than 10 for military vessels. Some of the extracted target features, such as orientation and centroid coordinates, are not used for automatic discrimination purposes but are instead are included as ancillary information in the final detection report. Extracting target measurements is a very simple

approach to target discrimination — it assumes the target has been delineated correctly and ignores the background context completely.

Interestingly, one of the most advanced ship detection systems, the Ocean Monitoring Workstation (OMW), only includes only a few such simple target measurements for target discrimination. Instead, the OMW relies on human supervision to identify and remove false alarms. Rey *et al.* give a good description of how this works in [102]. Following pre-screening, neighbouring detected pixels are clustered into a single detection. Detections which are too large, too small, too close to land, or too close to another detection are eliminated. All of these parameters are specified by the user. (Note that eliminated detections are not necessarily false alarms, they may simply be of no interest.) A maximum number of detections to display/retain is then specified. The detections are ordered according to size and presented to the human analyst in order. The human classifies each detection as a ship, a man-made object, or unknown; the presence of wakes is also noted. Rey *et al.* comment that the unknown targets are generally assumed to be sea spikes or other ocean structure but without ground truth the possibility that some are small boats remains. A similar philosophy of human supervision to remove false alarms seems to under-pin the system described by Cusano *et al.* in [25]. That system has no false alarm discrimination at all and relies on the accuracy of the prescreener. However, it does present a detailed display of each detection to the operator and includes wake detection and analysis at the operator's discretion.

In [49], Henschel *et al.* report on the utility of the OMW approach in dealing with false alarms. They note that while OMW appears to have an unacceptably high false alarm rate, the person-machine mix was found to work well in practice. To quote from the paper: *The false alarm rate is still an open question, however, our operational experience has shown that the person-machine mix, with very fast algorithms means that ships with specific type and classification characteristics can be supplied in a very short period of time. This means that in an operational scenario, the automated algorithm's false alarm rate can be ignored.* However, it is clear that there is still interest in developing a more automated approach to false alarm discrimination in the OMW. It was noted by Vachon *et al.* in [126] that this was one of the research directions at DREO. Yeremy *et al.* also discuss this issue in [142]. They note that small islands, buoys and ocean based oil rigs can be mistaken for ships by automatic algorithms such as the OMW but add that islands (with the exception of very small islands) and oil rigs can often be distinguished from ships visually. They imply that not much more can be done with single channel data but anticipate better progress with polarimetric data. As a pertinent example they note (see also [143]) that false detections caused by ocean buoys could be discriminated using the Cameron decomposition. As discussed in the section above on the Cameron decomposition, this is recent work and has not been developed into an algorithm yet.

We return now to a more detailed report of the target discriminators based on simple target measurements in the literature. We have already noted that the OMW rejects detections which are too large, too small, too close to land, too close to another target, [102]. The setting of these parameters is not discussed, presumably because they are scenario dependent. Wackerman *et al.*, [135], place a lower bound on the total target intensity. This is done mostly to avoid false alarms caused by noise spikes in areas with a high percentage of noise floor samples. But Wackerman *et al.* note that if a ship return needs to have energy larger than some threshold then this should be true anywhere in the

image. Robertson *et al.*, [105], likewise specify an absolute floor threshold . Ferrara and Torre, [35], place bounds on the ship area (number of target pixels) and the ship length and the width. A minimum distance between ships is also enforced. They do not specify the size of these bounds but note that they depend on image resolution and the size of the ships of interest. Casasent *et al.*, [19], work with SIR-C L-band SAR with a resolution of about 50 m. They consider detections with less than 5 pixels to be false alarms and remove them. They comment that this constraint removed all false alarms due to speckle noise.

Bounds on target area (size) can also be implemented using morphological filtering. Lin *et al.*, [68], work with ERS imagery with a resolution of about 12.5 m. False ship detections are eliminated using a morphological erosion filter. The filter operation is simple. It uses a 3x3 pixel moving window. If less than 3 of the 8 neighbouring pixels are possible ship pixels then the centre pixel is flagged as a false ship pixel and the detection deleted. In this way isolated false ship pixel due to speckle etc. are deleted. A different approach but with the same aim was used in an earlier paper, [67]. It is less convincing and we omit the details. Jiang *et al.*, [54] and [55], also remove false ship pixels with a morphological filter. They use the same erosion filter as Lin *et al.* but with a 7x7 pixel window and true detections require more than 7 of the 48 pixels to be detected. Liu *et al.*, [69], propose the use of two morphological filters: dilation and erosion. Dilation connects the neighbouring ship pixels and thereby clusters them. Erosion eliminates isolated detected pixels which are likely false alarms. No details of the filters are given.

Rye *et al.*, [107], discriminate false alarms on the basis that bright backscatter from a ship will occupy several pixels. Thus they suggest using a M of N discriminator whereby a ship is only detected if N pixels are above the detection threshold in an area of M pixels. They note that this can be further improved by taking pixel adjacency into account. As an example they proffer the use of two thresholds with a ship being detected if sufficient pixels are above the highest threshold and they are linked by pixels above the lower threshold. The system described by Rye *et al.* also includes wake detection. They note that if the location and extent of ship wakes in the imagery is known then false ship detections due to bright points on wakes can easily be discriminated. However, this requires more than just a limited search for wakes in the vicinity of the ship detections as is often done and may be computationally prohibitive in a near real-time system. They do not document the performance of his system — the aim was to apply it to ERS-1 data once the satellite was launched.

While the main thrust of the work by Gower and Skey, [43], is the general assessment of Radarsat ScanSAR imagery (50 m pixel spacing 100 m resolution) for fisheries monitoring rather than developing automatic ship detection algorithms, they do include an interesting feature extraction technique and associated target discrimination algorithm. They detect potential ships with a threshold and then fit a Gaussian "blob" model to extract discrimination features. The model is circularly symmetric with a profile of the form $A \exp(-r^2/a)$. The amplitude, A, width, a, and location are all derived by fitting the "blob" to the data. Thus in effect Gower and Skey are using a model-based approach to taking target measurements. Criteria are applied to the amplitude, A, width, a, and signal-to-background ratio, A/x, (where x is the rms scatter of nearby clutter). The fitting algorithm and detection criteria are not described. It was found by comparison with "ground-truthed" data that the detection criteria do allow some false alarms. No estimate

of missed detections is given.

Wackerman *et al.*, [135], identify two special types of false alarms: those in regions with a large number of system noise floor samples and those in land regions. Regions with a high proportion of noise floor samples can also contain bright spots which look like ship targets. They deal with these false alarms in two ways. First, a count is kept of the number of noise floor samples in the background window and when a high number are present, the prescreener detection threshold is raised. (This is done during the prescreening stage but it can also be done in the discrimination stage by reprocessing the data.) Second, as mentioned above, a lower bound is applied to the energy of the detect pixel (or pixels). This lower bound was set empirically by Wackerman using manually ground-truthed data. The other type of false alarm is the result of georegistration errors which lead to inaccurate land masking. In this case, Wackerman *et al.* identify a buffer zone where false alarms due to land are likely. In this zone they apply two constraints: the background mean and standard deviation must both be below thresholds. The reasoning here is that land is usually brighter and more variable than the ocean. The threshold on the background standard deviation is set empirically using manually ground-truthed imagery. The threshold on the background mean is more complicated — it is set dynamically for each image. Given the special nature of this latter threshold, we omit the details here and refer the reader to the original paper.

Eldhuset documents a complete system with extensive target discrimination in [32] and [31]. We review the target discrimination algorithms in the next few paragraphs. Following presceener detection, neighbouring detected pixels are clustered. Then a "ship-to-sea-condition" test is applied to discriminate false alarms. This test calculates the ratio of the integrated (or total) intensity of the ship to the contrast of the surrounding background and applies a lower bound. Eldhuset comments that the prescreener performs well in homogeneous seas but that the "ship-to-sea-condition" test is very necessary when swells, strong waves, eddies and steep fronts are present and that it reduces the number of false alarms substantially. Unfortunately he provides no details as to what the "contrast" of the background is or how it is calculated nor does he discuss how the value of the lower bound is set (he only states that it is high). Based on Eldhuset's earlier work, [30], it is likely that by contrast, Eldhuset means the normalised variance, σ_b/μ_b. Eldhuset states that the "ship-to-sea-condition" test is very rough and just a first step in reducing false alarms. The ensuing tests are more sophisticated but they rely on a wake being detected.

Following the "ship-to-sea-condition" test Eldhuset's system searches for a wake in the vicinity of the detection. A full description of Eldhuset's wake detection algorithm is beyond the scope of this report. Suffice it to say that the wake search done by averaging pixel intensities along scan lines to produce a set of scan curves which are then thresholded. If no wake is detected, the current detection is retained and processing passes to the next. If a wake is detected then the "homogeneity" and "scan curve" tests are applied. Eldhuset notes here that fronts, eddies and internal waves can confuse the ship and wake detectors. The scan curve test checks the behaviour of the scan curve around the potential wake feature. Eldhuset provides scant detail on how this is done and we quote from the paper *If an ordinary wake is present, the intensities on both sides of the wake are roughly equal, hence, the scan curve test rejects the wake if it (the scan curve) is too steep around the potential wake (feature)....* The idea of the homogeneity test is that the ocean area in front of a ship must be relatively homogeneous. Hence it is performed in a small sector C

in front of the ship. Homogeneity is tested by looking at the n-th central moments $\hat{\sigma}^n$ as estimated in the region C for $n = 2$ and $n = 3$. The region C is considered homogeneous if

$$E[\hat{\sigma}^n] - r\sqrt{\text{var}[\hat{\sigma}^n]} \leq \hat{\sigma}^n \leq E[\hat{\sigma}^n] - r\sqrt{\text{var}[\hat{\sigma}^n]}. \tag{59}$$

where $E[\cdot]$ and $\text{var}[\cdot]$ are the expectation and variance operators. Eldhuset details how $E[\hat{\sigma}^n]$ and $\text{var}[\hat{\sigma}^n]$ can be estimated using only knowledge of the mean image intensity μ (estimated in C) and the number of looks. Eldhuset states that for $r = 3$ the region is homogeneous with confidence 99.7% but in practice a larger value of the design parameter r is used. In [31], Eldhuset notes that this test is fast.

A caveat is placed on the use of the scan curve and homogeneity tests. While a ship and its wake can be rejected by a failure of either test, the homogeneity test is not applied to detections which are above a certain (unspecified) size. Such detections are retained with or without the wake depending on the outcome of the scan curve test. Several other items in the Eldhuset paper are of interest. First, the orientation of a ship is estimated by computing the moment of inertia about lines through its centre of mass. The line with least moment of inertia is taken as the ship orientation (with a 180^o ambiguity). For significant ships (more than 5 pixels) this orientation can be compared with the wake orientation (if a wake is present).

Eldhuset tested the system on Seasat imagery and ERS-1 3-look fast delivery imagery. The ERS-1 imagery was 3-look with a pixel spacing of about of 63 m azimuth x 50 m range. The imagery was visually ground-truthed and in total there were several hundred ships and wakes. Eldhuset reports that no false alarms were generated. A small percentage of ships were missed (6.8% in the Seasat data and 8.0% in the ERS data). A somewhat higher percentage of wakes were missed. Eldhuset adds that the missed detections can be reduced by fine tuning the various parameters but then the number of false alarms rapidly increases in rough seas. Lost and false wakes are more of a problem in the ERS imagery than the Seasat imagery.

We have already noted that the system presented by Ferrara and Torre in [35] uses elementary target measurements to discriminate false alarms. It also includes a wake detection module. Their aim is the detection of moving ships and so the presence of a wake is essential to the detection. The search for wakes is limited to the areas surrounding the ship candidates. It is based on an optimised version of the Hough transform. Rules ensuring minimum wake length etc are applied. Consideration is given to the different possible wake components. Wake direction and distance from the ship are estimated. Ferrara and Torre note here that ship speed can be estimated from this data as long as the satellite velocity is known too. In the final section of the paper, Ferrara and Torre discuss how an ISAR technique can be used to construct 2D or 3D images of the moving targets. Obviously such images would be useful for classification purposes. In [34], Ferrara and Torre report on the test phase of the system where the optimal parameters for the various rules have been determined and system performance has been measured. Unfortunately, no details of the settings are given. Also, while good results on real data are reported, their reliability is questionable since there is no discussion of the false alarm rate and the false alarms appear have been included in the estimation of the probability of detection.

In subsequent work, [33], Ferrara *et al.* propose to improve the detection of moving ships by using a moving target detector based on local residual coherent phase error in

parallel with the wake detector. The target motion with a component in the azimuth direction manifests itself locally as a coherent phase error. This phase error has a strong quadratic component in the azimuth direction which directly depends on the azimuth velocity component. The motion detection technique is derived from the Phase Gradient Autofocusing algorithm and allows a robust estimation of the gradient of the coherent component of the phase error. The algorithm is iterative with rapid convergence. An outline is given in the paper and references with more detail are cited. While computationally expensive, it need only be applied locally in the vicinity of the detected ship pixels. Any motion detected in this way corroborates the presence of a ship. An advantage of this approach over wake detection is that wakes are not always imaged whereas the motion effect is. The method was tested on SIR-C/X-SAR SLC data. The results are presented as plots of phase against frequency samples. From this it was concluded that ship forward motion could be easily detected and that roll, pitch and yaw motion effects could also be seen to a lesser extent. This work was investigative and no suggestion was made as to how an algorithm (with thresholds etc.) could be developed from it.

Techniques for false alarm identification and removal for the SUMO detector are given by Schwartz *et al.* in [114]. While we have already reviewed that paper in the section on the SUMO detector, it is appropriate to provide a brief summary here too. Schwartz *et al.* find there are two main types of false alarms produced by the SUMO detector: bright ship-like sidelobe patterns on large objects and bright ship-like noise from local wind turbulence effects. Dealing with the first type is easy. Any detected pixels which are less than 8 pixels apart are grouped together. Thus false alarms due to sidelobes are incorporated into the detection of the main object. Dealing with the second type of false alarm is a little more complicated and depends on the SUMO detector thresholds. Essentially, the discriminating feature is a large number of bright local maxima in the vicinity of the detection. For more details see the description in the section on the SUMO detector or the original paper. Other techniques for false alarm reduction in the SUMO detector have also been considered. This work is outlined in a presentation by M. Alvarez to the inaugural IMPAST meeting in January 2002. the presentation is available from the the IMPAST web site, [57]. The idea is to extract contour and shape information from the detections. A decision tree for differentiating fishing boats from oil rigs is proposed. No further mention of this work has been found and we are uncertain as to its effectiveness or current status.

It has already been noted that the system Lin *et al.* work with, [68], uses a threshold to identify potential ship pixels and then an erosion operation to remove false alarms. Following this neighbouring ship pixels are clustered into a single detection and ship orientation is estimated using a total least squared error regression. A wake detection phase is then employed in the vicinity of each potential ship detection. The wake detector uses a Radon transform technique based on that of Copeland but with major modifications. False wake components are eliminated using three different wake tests. One of these tests compares the wake orientation with that estimated from the ship. Wakes were found to be much harder to detect than ships. In light of this, we assume that the presence of a wake is not considered necessary for ship detection but no comment is made on this point in the paper.

Casasent *et al.*, [19], note that false alarms can occur in the sidelobe patterns of ships with large sidelobes. They handle this case by using a filter to detect the sidelobes and

then mask out any detections within them. The filter was constructed from real data by forming an "average" sidelobe pattern. Details are in the paper. In the final section of the paper, they also propose a set of rules aimed at reducing the false alarm rate when the background is very noisy, very dim or very smooth. The rules are *ad hoc* and mostly involve adjusting the value of $\sigma_b t$ in (15) based on the size of μ_b and σ_b. We omit the details.

Robertson *et al.*, [105], found that the CFAR detector they used produced an unacceptably high FAR. Consequently, a neural network algorithm was applied as a post-processor to filter out false alarms. The neural network is a Kohonen neural network and requires training on a representative ground-truthed image. A reference to work at DERA Malvern on the neural network is given. A final detection performance of 88% probability of detection against a FAR of 12% is given for Radarsat ScanSAR imagery. Curiously, the FAR seems to be calculated as a percentage of all ships visible in the image which is highly unusual.

In a pair of related papers, Benelli *et al.*, [8], and Argenti *et al*, [5], describe a method of discriminating targets which is based on a fuzzy classifier. The input is a list of potential "ship-blobs", i.e. patches of bright pixels which the prescreener has identified as potential ships. Three features are extracted from these blobs for use in the classifier: area, mean grey level and elongation (we presume length is meant here). The classifier uses a "Weighted Distance Method" to calculate the membership function $\mu_C(x)$ where x is the blob being processed and C is the set of all true ships. A high value of $\mu_C(x)$ (i.e. close to 1) indicates the classifier believes the blob to be very ship-like while a low value (i.e. close to 0) indicates the opposite. The calculation of $\mu_C(x)$ is based on the weighted Euclidean distance in feature space of x from the various prototype ship blobs. We omit the details here. We do however note that the weights depend on the importance assigned to each feature. Benelli *et al.* have set the importance of the features, area, mean grey level and elongation, to be 10, 6 and 5, respectively. They also state that the prototypes were manually selected by human experts during a training phase. The technique was tested on Seasat SAR imagery but no detailed performance figures are given.

Finally, we mention the work being undertaken at the Universitat Politecnica de Catalunya (UPC) by Mallorqui *et al.*, [80], [79], [78] and [77]. This work involves the simulation of SAR imagery of ships and is part of the IMPAST program. IMPAST is aimed at improving fisheries monitoring using SAR imagery. We discuss IMPAST in the next section. Margarit *et al.*, [80], state that one of their aims is to provide a tool for vessel classification studies. They note that large databases containing ground-truthed SAR imagery of ships for wide range of different ships with different orientations, radar parameters and ocean conditions are needed to train classification algorithms and that such data are not available. They propose that the simulator be used to produce training data for classification algorithms. They add that it can also be used to optimise SAR system parameters for ship detection — it should allow the effects of incidence angle, ship orientation, etc. to be studied. The simulator uses realistic ship models and takes account of vessel velocity and vessel motion due to ocean swell. It is also fully polarimetric. At present the simulator is still under development and the associated classification studies have not begun.

9 Ship detection systems and research groups

In this section, we review the various ship detection systems which have been documented in the literature and outline how they work as a whole. We also discuss the research groups and people involved in the development of these systems. We do not discuss the component algorithms in any detail here since that has already been done in earlier sections. We include indications of possible future developments.

Several near real-time operational ship monitoring systems using SAR imagery are now in place or are currently being developed. The prime examples are the Ocean Monitoring Workstation (OMW), [126] and [47], the Alaska SAR Demonstration (AKDEMO) system, [135] and [94], the European Community Joint Research Centre (JRC) system, [114] and [61], and the Qinetiq's MaST system, [96]. The Norwegian Defence Research Establishment (NDRE) is an important player too and was behind the system described by Eldhuset in [32]. Other examples include the work of Robertson *et al.*, [105], Cusano *et al.*, [25], Lombardo *et al.*, [70], Ferrara *et al.*, [33] and Lin *et al.*, [68]. We discuss each of these systems in turn.

The OMW was developed by Satlantic Inc, with financial and technical input from the following organisations: Canadian Centre for Remote Sensing (CCRS), Canadian Department of Fisheries and Oceans, Canadian Coast Guard, Canadian Department of National Defence and Canadian Space Agency, [134] and [133]. The OMW was designed to provide near real-time analysis of Radarsat imagery (although it can use ERS and other SAR imagery). It grew out of the Ocean Features Workstation, see [103] and [48], and its development has been ongoing ever since. In addition to ship detection, the OMW can extract wind vectors, calculate 2-D wave spectra, perform oil slick detection, and classify other ocean features. Vachon *et al.*, [126], note that the Canadian Department of Fisheries and Oceans has been evaluating the operational utility of the OMW for fisheries monitoring and that the Canadian Department of National Defence, through the Defence Research Establishment Ottawa (DREO) was also evaluating the OMW for offshore ship surveillance.

General descriptions of ship detection with the OMW are given in [47] and [134] and a more detailed description is given in [102]. It begins with an image ingestion phase. In this phase, data preprocessing, georeferencing, land masking and beam seam/nadir ambiguity masking are undertaken. The World Vector Shoreline database is used for land masking. The next phase is the prescreening phase. The prescreener is a standard CFAR detector based on fitting the surrounding ocean with an L-look K-distribution probability density function. Rather than using a moving window, the CFAR threshold is calculated for and applied to whole images tiles at a time. Details on this detector are given in [49], [128] and [14]. Following that, is a discrimination phase with two parts. The first part is automatic and is based on simple target measurements such as size and proximity to other targets and land. The second part is manual and involves presenting the detections to a human operator for confirmation. The final OMW output is a report on candidate ship positions and estimated sizes. According to Vachon *et al.*, [134], if a wake is present, the report also includes ship heading and speed based on analysis of the wake. Many technical details on the implementation of the OMW are in the OMW manual, [111], which was available

from Satlantic's web page [9]. Many of the algorithms are user configurable. While dealing with the false alarm rate using the final person-machine mix is not entirely satisfactory, it has been shown to work well in an operational context, see [134], [49] and [47].

It appears that much of the current/future research effort on the OMW is/will be aimed at reducing the false alarm rate (FAR) through the use of polarimetric data. In [126], Vachon *et al.* note that DREO's research efforts at the time were to reduce the FAR of the OMW and modify the OMW to make use of fully polarimetric SAR data. In their report on the AMRS ship detection workshop in June 2000, Vachon and Olsen, [132], note that J. Campbell (DREO, Canada) discussed a ship detection validation exercise using airborne C-band polarimetric data and he indicated several ways polarimetry may improve detection performance. Vachon and Olsen also note that M. Henschel (Satlantic Inc, Canada) discussed the OMW, including plans to upgrade the OMW so that it can use polarimetric data and also fuse targets detected in SAR with other data from other sources. Yeremy *et al.* also discuss the false alarm rate issue in [142]. They imply that not much more can be done with single channel data but anticipates better progress with polarimetric data. As a pertinent example they note (see also [143]) that false detections caused by ocean buoys could be discriminated using the Cameron decomposition.

The ship detection system described by Wackerman *et al.*, [135], was part of the American NOAA/NESDIS (National Oceanic and Atmospheric Administration/National Environmental Satellite, Data and Information Service) program called the Alaska SAR Demonstration (AKDEMO). The goal of the Alaska SAR Demonstration was the demonstration of near real-time Radarsat-1 SAR applications in Alaska. Of course, given the long coastline and large continental shelf in Alaska, ship detection and fisheries monitoring in particular, is one of the more valuable applications. More information on the Alaska SAR Demonstration and the ship detection system may be found in the paper by Pichel and Clemente-Colon, [94]. Related work validating the system is documented by Friedman *et al.* in [38] and [39]. It is noted by Friedman *et al.* that the CFAR detector the system is based on was developed by Veridian ERIM International. However, we will refer to the system as the AKDEMO system.

Like the OMW, the AKDEMO ship detection system is fairly simple at heart. It has an initial land masking step which includes a ± 2 kilometre buffer zone of possible land masking errors. The buffer zone requires special processing in order to avoid a high false alarm rate. The ship detection is performed with a moving window implementation of the two parameter CFAR detector as discussed in Section 5.3. No general false alarm discrimination is undertaken. However, two special cases of false alarms are dealt with. The first occurs in regions with a high proportion of SAR system noise floor samples and second occurs in the land masking buffer zone. Details of how these are dealt with may be found in Section 8 or in the original paper, [135]. Validation of the system is discussed in, [38] and [39] as well as in [135] and good performance is shown. One scant reference to possible future developments was found. Pichel and Clemente-Colon, [94], indicate that Envisat data will be input into the AKDEMO. Presumably some research will be directed towards making use of polarimetric data for ship detection.

The fisheries control group at the Joint Research Centre (JRC) of the European Community, are developing a ship detection system for monitoring fishing vessel activity. This

---

[9]http://www.satlantic.com

system, called SUMO, has been described in several papers, [114], [61] and [60]. More information can be found on the JRC web site, [57]. See also the paper by Wagner, [136] and the article by Kourti and Shepard on the RSI web site, [58]. The development of SUMO is currently very active and the following description comes mostly from [61] and so may not be completely up to date. Unfortunately, more recent documentation could not be found. As usual, the first step is land masking. Originally this was done manually but it now uses the publicly available Global Self-consistent Hierarchical High-resolution Shoreline (GSHHS) database. The land mask is not always accurate and a buffer zone is recommended. Following land masking the image is partitioned into tiles and the template based two parameter CFAR detector described in Section 5.7 is applied to each tile. Detected pixels are then clustered and detections which are too large to be fishing vessels are removed. The final step is to apply the SUMO discrimination algorithms described in Section 5.7. The system is written in Java. Validation of the system was reported on to some extent in [114] and [61]. The system has been used in many fishery monitoring campaigns and further performance information is available from the JRC web site, [57].

One of the least convincing aspects of the SUMO ship detection system is the setting of the threshold factors in 30 (and 31). This was done by trial and error. While the factors in 30 are a good overall compromise, Kourti *et al.*, [61], found that the best choice of factors varied from region to region. Thus one of the future goals for SUMO is a method of automatically selecting the best combination of threshold factors. Another future research direction is improved discrimination/classification of the detections. It is noted on the JRC web site, [57], that a database of all identified targets is being built for this purpose. The database is reported to have about 20 vessels with more than 50 radar cross sections and about 20 oil rigs with more than 60 radar cross sections. The aim is to discriminate fishing vessels from oil rigs and other ships and further to discriminate different classes of fishing vessels. A third planned improvement of the system is the ability to process Envisat data. While this is unlikely to make much difference in terms of the detection algorithms for single channel data, Envisat includes two channel data in the Alternating Polarisation mode and it is likely that changes will be needed to make best use of it.

Wagner, [136], reports in an interview with Kourti, that one of the outcomes of JRC's research into automatic detection of ships in SAR imagery was the setting up of a program called **IMPAST**. IMPAST is an "acronym" for Improving Fisheries Monitoring Through Integrating Passive and Active Satellite Based Technologies. More information on IMPAST can be found on the JRC web site, [57]. See also the article, [58], on the RSI web site. The aim of IMPAST is to improve fisheries monitoring by combining the detection of ships in satellite SAR imagery with ship position reports from the Vessel Monitoring System (VMS). The VMS is not explained in detail on the IMPAST web site but briefly the idea is that participating ships report their geographic location (and other information) when requested. Participation in the VMS is a requirement of fishing vessels in certain fisheries. By correlating the VMS reports from a fishery, with ship detections in SAR imagery of the fishery, it is possible to identify vessels in the imagery without associated VMS reports and therefore possibly in breech of fishing regulations. One of the specific aims of IMPAST is to develop a near real-time pre-operational automatic system for doing this. While the development of good ship detection algorithms is an essential part of this system, much effort has also gone into other parts of the system, such as data management

and correlating the ship detection with VMS reports. Note that one of the challenges in IMPAST is sorting out non-fishing ships from fishing boats. Note also that the VMS in effect provides partial groundtruth for the SAR imagery and hence can help develop better vessel detection tools. We add, however, that VMS is only required for fishing boats longer than 24 m although there are plans to have this reduced to 18 m and then 15 m over the next two years. Further, VMS non-reporting and mis-reporting are obvious problems as is synchronising the VMS reports with the SAR data collection.

IMPAST is a three year project which began in January 2002. It involves 14 research and commercial partners. Six of these partners are European fisheries monitoring centres and hence represent the end users. The development of SAR image analysis tools is being handled by the 3 scientific partners: JRC, Qinetiq and UPC. The JRC is developing a ship detection system called SUMO while Qinetiq is developing one called MaST. We have described SUMO above and we will describe MaST below. The UPC is working on the simulation of SAR imagery of ships. Amongst other things, the UPC work will lead to the generation of training data for vessel classification algorithms and the development of the algorithms themselves. Radarsat is the main source of imagery for IMPAST but Envisat is also being considered. Several IMPAST progress meetings have already been held. Some of the material presented at these meetings and related documents are freely available from the JRC web site, [57], but some is also not publicly available. The most recent meeting was held in February 2004 with one of the significant items being reports on the monitoring of three fisheries in mid-2003. This was obviously a good opportunity for validating the various systems. A presentation to the EU Fisheries Expert Group meeting in June 2003 is also available from the JRC web site. It lists the 2003 aims of IMPAST as: set up and test the whole system; add Envisat imagery; improve the fusion of VMS reports with detected ships; improve the discrimination/classification algorithm; and involve more member states. Full operational testing and a cost-benefit analysis is planned for 2004.

A closely related project is DECLIMS. This project is relatively new, with the inaugural meeting in May 2003, and is yet to mature. Currently, the main source of information on DECLIMS is the web site, [57]. We quote from the web site: *The project DECLIMS, Detection and Classification of Marine Traffic from Space, was defined in order to provide a focus for research into the use of satellite imagery for maritime vessel detection, classification and identification so as to allow researchers to develop a better understanding of the capabilities of such systems, to identify the advantages and drawbacks of different approaches, to strengthen the infrastructure capable of meeting the demands of users, and to help drive developments of new sensors and platforms towards the operational need of vessel monitoring.* The project involves 16 partners, led by the EC-JRC and including Qinetiq, CCRS, Radarsat International, TNO and FFI. It is evident from the web site that among other things, DECLIMS aims to: provide a benchmark for SAR ship detection; provide a benchmark for SAR ship classification; and provide a benchmark for wake detection. As yet however, there is no publicly available output from DECLIMS.

Qinetiq has developed a ship detection system called MaST. While it is clear from the JRC web page, [57], that MaST is a well developed product, there is very little documentation available publicly on its component algorithms. Presumably this is because it is a commercial product. The following incomplete description of MaST was gleaned from the demonstration web page, [96]. The basic steps are: land masking; detection of individual

ship pixels; clustering to remove false alarms; and calculation of ship measurements. No details on the detection algorithm are available although mention is made of a CFAR detector. The calculated ship measurements are location, area, perimeter, length, width and direction. A graphical user interface is included which allows the user to select targets and review and categorise them. Given Qinetiq's involvement with IMPAST and the commercial nature of the product, it is presumed that the system detection performance is satisfactory if not better. The JRC web page, [57], contains several Qinetiq presentations to the IMPAST progress meetings and these include reports on system performance. In the second progress meeting in January 2003, an overall probability of detection of 73% was reported but no mention was made of the false alarm rate. As well as detecting ships in SAR imagery, MaST can extract ice, wave and oil slick information. Future plans include the detection and analysis of other ocean features.

The Norwegian Defence Research Establishment (NDRE) has long been interested in ship detection systems for SAR imagery. The system described by Eldhuset in [32] was developed at NDRE. That system was based on earlier research for NDRE by Aksnes, Wahl, and other. See for instance the paper by Wahl *et al.*, [138], which reports on an European Space Agency sponsored project titled "Ship Traffic Monitoring using the ERS-1 SAR" and an earlier one, [137], which reports on an ESA contract titled "SAR-Detection of Ships and Ship Wakes". Eldhuset notes that the system has been commercialised and was sold to Mitsubishi in Japan. While Eldhuset's work was published in 1996, it is clear that NDRE has an ongoing interest in ship detection. In the paper, [89], Olsen and Wahl report that Norway has established an operational service for fisheries monitoring which makes significant use of SAR imagery. In their summary of an AMRS Workshop in June 2000, Vachon and Olsen, [132], note that R. Olsen (FFI, Norway) discussed ship surveillance activities in Norway and stated that this has now reached operational status and development is ongoing. More recently, Olsen and Wahl, [90] have assessed the ship detection capability of Envisat SAR data. Note too that Olsen is a co-author of several papers on the OMW. Finally, as will be mentioned at the end of this section, Olsen has a current ENVISAT-AO titled "Ship detection with Envisat using alternating polarizations". However, Eldhuset's paper, [32], remains the most recent description of a Norwegian ship detection system.

The Eldhuset system, [32], is composed of five modules. In the first two modules, the SAR image is georeferenced and land areas are masked out. In the third module, ships are detected with what is effectively a cell-averaging CFAR detector and then false alarms are removed with a "ship-to-sea condition" test. The "ship-to-sea condition" test is very necessary since the CFAR detector threshold is set at a low value. Details of these algorithms are discussed in Sections 5.3 and 8. Wake detection and analysis are undertaken in the fourth module. The final module relies on both the ship and wake detection results and is aimed at further false alarm reduction. Again, we have discussed the details in Section 8. Eldhuset has tested the system extensively and reports that it works well on both Seasat and ERS data. This system was one of the first comprehensive and well documented systems and is often cited in the literature. While no recent reports on the system have been found in the literature, it is clear that Norway through the NDRE has an ongoing interest in ship detection systems. In their paper, [89], Olsen and Wahl report that Norway has established an operational service for fisheries monitoring which makes significant use of SAR imagery. Also, as mentioned below, both Olsen and Wahl have

current ENVISAT-AO proposals. Finally, NDRE is one of the participants in DECLIMS.

Robertson, from the then DERA Space, U.K., presented a ship detection system at the Ship Detection in Coastal Waters Workshop organised by Alliance for Marine Remote Sensing (AMRS) in June 2000, see [132] and [105]. The latter reference provides the only details we have on this system. The first step is geo-referencing and land masking. The system then applies a standard CFAR detector based on modelling the ocean clutter with a K-distribution. However, unlike the OMW detector, this detector also uses a standard moving window setup like that in Figure 1. An absolute floor threshold is used to avoid false alarms in dark areas of the image. The false alarm rate is then further reduced by applying a Kohonen neural network classifier. The system has a graphical user interface which allows imagery and detection parameters to be specified and then displays the detection results. The system was validated using Radarsat ScanSAR Narrow Far imagery of the English channel.

Several papers on ship detection have been published by researchers in the Department of INFOCOM - University of Rome "La Sapienza", Cusano *et al.*, [25], Lombardo and Sciotti, [70] and [115], and Sciotti, Pastina and Lombardo, [117] and [116]. The paper by Cusano *et al.*, [25], describes a real-time operational ship detection system with application to ERS-1/2 imagery. The system is installed in the ESA-ESRIN site and processes "quick look" imagery. The system is composed of four modules: land masking, ship detection, ship wake detection, ship classification and display of the results. The first two modules are applied to the whole image while the last two modules are only applied to operator selected targets. The ship detector, is based on a multi-target pixel, cell averaging CFAR detector (16) but it has been modified in order to cope with the high variability of sea surface scattering. In addition, areas of very high contrast are automatically detected and masked out. Wake detection is based on Eldhuset's scan method, [32], but with some improvements. Ship velocity is estimated in the usual way from the displacement between the ship and its wake. The ship classification is coarse, using only size and speed. The results show that the detector works well and is not confounded by closely spaced ships or the presence of ship wakes. The paper lists the following future research directions: i) improving the ship detector by using context-based analysis to cooperate with CFAR schemes 2) improving the wake detection and analysis module 3) comparing the present technique with a segmentation based approach.

It appears that one of the future research directions identified by Cusano *et al.* was followed up in the work of Lombardo and Sciotti, [70] and [115]. Lombardo and Sciotti propose making ship detection more context based by including a preliminary segmentation stage. The idea is to partition the image into regions with similar backscattering behaviour and apply different detectors in each region. In this way the high false alarm rate produced by CFAR detectors in inhomogeneous imagery can be overcome. We have reviewed Lombardo and Sciotti's work in detail in Section 5.5. We only comment here that Lombardo and Sciotti's results are convincing and we believe this is an important approach to the ship detection problem. Lombardo and Sciotti note that prior segmentation should be a good basis for other information extraction algorithms, such as ship wake detection, pollution detection and wind retrieval. Hence, this is one of their future research directions. They also indicated their intention to investigate the segmentation of polarimetric SAR data and whether it could help improve detection performance. The latter was done by Sciotti, Pastina and Lombardo in [117] and again it was demonstrated

that segmentation does improve detection performance.

Ferrara *et al.*, [33], [34] and [35], describe a ship detection system developed by the Italian company Alenia Aerospazio. This system grew out of a more general program on SAR data analysis and automatic target recognition at Alenia Aerospazio which began in 1994. Ferrara and Torre outline the basic system in [35]. It is aimed at detecting moving ships and so ship wake detection is considered to be an essential step. The initial steps consist of speckle filtering, coastline detection and land masking. Following that potential ship pixels are detected using a CFAR threshold method. The ship-like pixels are then clustered and length and width are calculated. Detections are subject to constraints on size, length, width and proximity to other detections. The remaining detections are then passed to the wake detection module. Wake detection is based on the Hough Transform. Wakes must satisfy a minimum length. As noted, detections without a wake are rejected. In the final step, ship heading and speed are estimated. Ferrara and Torre discuss the performance of the system in [34]. However, this is not a particularly convincing paper — apart from an interesting theoretical analysis of the probability of detection, the results are poorly presented and explained. In 1998 Ferrara *et al.*, [33], discuss some improvements to the system. In particular, they: introduce an additional preprocessing step to suppress SAR sidelobes via spatially variant apodisation; improve the CFAR detector by allowing the automatic selection of a cell-averaging CFAR detector or an order statistics CFAR detector; and use a moving target detector based on local residual coherent phase error in parallel with the wake detector. Again, validation of the system performance is lacking. The current status of the Alenia Aerospazio ship detection system is uncertain since no further references to it have been found.

A ship and ship wake detection system has also been developed by Lin and other researchers, [68] and [67], at the Centre for Remote Imaging, Sensing and Processing (CRISP) in the National University of Singapore. Their aim was to use ERS SAR imagery for monitoring and routing of ship traffic around Singapore. They note that Singapore is one of the busiest ports in the world. Not much has been published recently about the CRISP system and according to the CRISP research web page the ship detection system is no longer an active area of research. The system consists of four distinct stages: a pre-detection stage where land masking and calibration to normalised RCS are undertaken; a ship detection stage where global thresholding followed by morphological filtering and clustering of detections are undertaken; a wake detection stage using the Radon transform; and final stage where ship heading and speed are estimated. There is no published record of system validation. Lin *et al.*, [68], indicate that further development is needed to make the system more robust to rough sea and wind conditions.

The future plans of some researchers can be anticipated by searching the Announcement of Opportunity proposals for the Category 1 scientific exploitation of ESA ERS and Envisat data. The idea of these proposals is to maximise the beneficial use of ESA data from both ERS and ENVISAT satellites and to stimulate a balanced development of science, public utility and commercial applications, consistent with the mission objectives. The accepted proposals are listed on the ESA web site [10]. The following entries can be found: ID 255, ENVISAT-AO, P.W. Vachon, Canadian Centre for Remote Sensing, "Evaluation of ship detection and classification by the Envisat advanced synthetic aperture radar

---

[10]http://eopi.esa.int

alternating polarisation mode"; ID 411, ENVISAT-AO, R.B. Olsen, Norwegian Defence Research Establishment, "Ship detection with Envisat using alternating polarisations"; ID 494, ENVISAT-AO, T. Wahl, Norwegian Defence Research Establishment, "Feasibility of Envisat ship detection for fisheries enforcement in the South Atlantic"; and ID 2261, Category-1, O.M. Olsen, Kongsberg Spacetec AS, Norway, "Ship detection and oil spill detection in MASS using ASAR Wide Swath Mode data." The equivalent scheme for Radarsat-2 data is the Science and Operational Applications Research (SOAR) program administered by Radarsat International and the Canadian Space Agency. Information on SOAR may be found on the Radarsat-2 web site [11]. However, proposals are still being solicited and the winning proposals will not be announced until November 2004.

10 Recommendations

The generic problem addressed in this report is the detection of ships in ocean areas of interest. However, it is not possible to do this perfectly and various tradeoffs are required. These tradeoffs are made within the context of surveillance scenarios with specific objectives. Examples of the considerations that need to be taken into account are:

- Size of the surveillance area;
- Type, class, size and velocity of ships to be detected;
- Ocean environment in the surveillance area — sea state, wind conditions, etc.;
- Time constraints and revisit rate;
- Available computing resources.

The radar parameters also need to be taken into account. The primary ones being:

- Frequency;
- Resolution;
- Incidence angle;
- Polarisation

Ideally, one would have a particular surveillance scenario in mind and the radar parameters would be chosen accordingly. Ship detection algorithms would then be designed to suit the resulting imagery. However, this report is not aimed at addressing any specific surveillance scenario. Rather it aims to be a general review of techniques for ship detection in SAR imagery and hence is seen as a starting point for tackling any particular surveillance problem.

Having indicated the broad scope of our outlook, we remind the reader that the important topic of ship wake detection is not addressed here. It is anticipated that wake

---

[11] http://www.radarsat2.info

detection will be the topic of a follow up report. The pros and cons of ship wake detection were discussed in Section 2.5. We also note that as a reflection of the bias in the ship detection literature towards space-based SAR imagery, airborne SAR imagery has not really been considered. Fingas, [37] and [36], points out that this should make little difference because the imaging principles for both platforms are the same. However, we note that space-based SAR imagery is generally single channel and lower resolution. Thus most ship detection algorithms in the literature are aimed at locating bright point-like targets in single channel imagery. Space-based SARs also tend to have lower incidence angles which makes the detection problem harder. We add however, that the next generation of space-based SARs offer multi-channel polarimetric data and consequently there is a shift towards research on algorithms which make use of multi-channel data.

One of the most basic decisions to be made is the resolution/coverage tradeoff. Radar design limitations mean that increased resolution generally comes at the expense of decreased swath width. Thus while higher resolution allows for higher probabilities of detection and especially for smaller ships, it comes at the cost of narrower swath widths and hence decreased coverage. Generally, the coarsest resolution which still allows good probabilities of detection is chosen, so that coverage is maximised. Thus Radarsat ScanSAR Narrow imagery (50 m resolution and 300 km swath width) is acceptable for the surveillance of ocean-going fishing vessels but higher resolution imagery such as Radarsat Standard (25 m resolution and 100 km swath width) or Fine (9 m resolution and 45 km swath width) may be needed for the surveillance of smaller vessels in coastal areas.

Any further choice in radar parameters should clearly be made to maximise ship detectability. We discussed ship detectability in Section 2.4. The main points are that: finer resolution is better; large incidence angles are better; HH polarisation is better than VV; and HV polarisation is better than HH at incidence angles less than about 45^o. We also remind the reader that SAR imagery is sensitive to surface winds and in severe conditions even large ships may not be visible. Likewise ship construction material is relevant and small wooden or fibreglass boats are often not visible.

10.1 Benchmark testing

A difficulty in making recommendations on ship detection algorithms is the lack of rigorous performance evaluation which makes comparison of the various approaches difficult. Some algorithms are very poorly tested with results only being presented for a few images. In such cases really all that has been demonstrated is the feasibility of the approach. For those algorithms which are more extensively tested, it is still often the case that the test data lacks rigorous groundtruth. It is very common for the groundtruth to be provided by visual inspection of the imagery. This makes it hard to accurately estimate probabilities of detection and false alarm rates. Further, different researchers use imagery from different SAR systems and ocean environments making comparisons even more difficult. While many researchers may have made their own comparisons, only a few researchers have clearly documented such comparisons. In other words, benchmark tests are severely lacking. As discussed in Section 9, provision of benchmarks is one of the aims of DECLIMS. **Given this discussion our first recommendation is to:**

1. **Specify a surveillance scenario.**

2. **Decide on appropriate radar parameters and geometry.**

3. **Set up a well ground-truthed benchmark test.**

4. **Compare the different detection algorithms.**

We note that procuring well ground-truthed imagery for ship detection purposes is likely to be time consuming and expensive and this is probably explains why it has been lacking to date.

One way of overcoming the lack of ground truth is to simulate data. As mentioned at the end of Section 8, this is the approach being taken by the researchers at the Universitat Politecnica de Catalunya (UPC). The UPC researchers have focused mainly on ship backscatter with the aim of providing a tool for ship classification studies. However, it is reasonable to expect that this type of approach could be extended to include the simulation of ocean backscatter for ship detection studies as well. **The advantage of simulation is that it would allow the generation of benchmark tests which include a wide variety of ocean conditions, radar parameters, ship types and ship orientations.** Of course there are limitations to simulations and the results would still need to be checked against real ground-truthed data.

10.2   A basic system

Our intention in the remainder of this section is to provide guidance to the selection and testing of algorithms. The question then is: which of the existing algorithms should be tested (and in what order)? We answer on the basis of:

- good performance results from well conducted experiments;

- ease of implementation and testing;

- experience;

- use of non-standard approaches and or features.

The rational behind the last item is that non-standard approaches are likely to use complementary information to the standard approaches and therefore when used in conjunction with standard approaches are likely to improve performance.

The first item is clearly the most important. In relation to this we note that the only complete systems to have extensive and convincing publicly available reports of their performance are the OMW, AKDEMO, SUMO and Eldhuset systems. **On this basis we recommend the following basic system for low to medium resolution single channel imagery:**

1. **Land masking using a shore-line database with a buffer zone included;**

2. **Detection with a simple moving window adaptive threshold algorithm;**

3. **Clustering of detected pixels;**

4. **Discrimination of false alarms using simple target measurements;**

5. **Human supervision and discrimination of the final detections.**

We note that this system can be applied to higher resolution imagery if an initial down sampling step is included. Once this basic system is in place, improvements and alternatives can be considered. We discuss the details of the basic system and make suggestions for improvements and alternatives in the remainder of this section. We believe that while many of the refinements and alternatives may result in improved performance, **the most significant gains are likely to come from the use of polarimetric data and from better false alarm discrimination.**

10.3   Land masking

Although false alarms from land are likely to be easily discriminated, some form of land masking is needed to prevent the sheer number of false alarms from overloading the system. **The recommended method of land masking is to use a shore-line database.** This is done in the OMW, AKDEMO, SUMO and Eldhuset systems as well as many others. It may be useful/necessary to improve masking accuracy by allowing the user to intervene in the masking process. In order to cope with masking errors it **is also usual to include a buffer zone around the mask edge.** Detections in the buffer zone can be either rejected or subject to special false alarm discrimination as is done in the AKDEMO system. We also note here the possibility that the azimuth shift of a fast moving boat near the shore-line may place its backscatter on land.

A desirable extension to land masking algorithms is the facility to implement user defined masks. Such masks may be needed for: correction of "errors" in the land mask; masking out tidal or sea ice areas; and testing system performance over specified areas in the imagery (where the ground truth is known).

Land masking with a shore-line database may not always be sufficiently accurate, and especially if the littoral zone is of interest. Tidal variations in the water-line, eroding coastlines and unmapped rocks and shoals may need to be taken into account. In this case, automatic land detection algorithms need to be considered. Only the Alenia Areospazio ship detection system, see [33], actually uses automatic land detection. That algorithm is based on the work of Lee *et al.*, [64]. However, there is an extensive literature on this topic which has not been reviewed in this report. **If automatic land masking is deemed to be necessary, we recommend that the literature be reviewed and the different approaches compared.**

10.4   Preprocessing

Most ship detection systems do not make use of special preprocessing algorithms. Of the few that do, the technique used is some form of speckle filtering. As we have noted before, speckle filtering will alter the image statistics and the ensuing detection algorithm must be compatible with this change. Since we will recommend the use of standard CFAR detectors which are likely to be incompatible with speckle filtering we do not recommend the use of speckle filtering as a general preprocessing technique.

10.5 Single channel detectors

In recommending detectors for single channel imagery we first note that the nature of the imagery needs to be taken into account to some extent. Different types of imagery require different detectors or at least different detector parameters for best results. For instance, Wackerman *et al.*, [135], use several different threshold factors depending on the resolution of the imagery. Lombardo and Sciotti, [70], use different segmentations for different resolution imagery. The JRC researchers, [4], have found that SUMO does not work well on logarithmically scaled data. On this latter point we note that the same will be true of the two parameter CFAR detector on which SUMO is based.

We have already noted that the OMW, AKDEMO, SUMO and Eldhuset systems are the most convincing ship detection systems described in the literature. All these systems use an adaptive threshold algorithm for prescreening. Further, the prescreeners used in the AKDEMO and SUMO systems are both based on the standard two parameter CFAR detector. This detector is CFAR for Gaussian distributed data but as Wackerman *et al.* have pointed out, [135], it is also CFAR for gamma and K distributions as long as the normalised variance is constant. **Hence the two parameter CFAR detector is an obvious starting point and it is our recommended first option.**

The AKDEMO system as described by Wackerman *et al.*, [135], uses the standard multi-target pixel two parameter CFAR detector given in (15) with the standard moving window setup shown in Figure 1. On the other hand, the SUMO detector uses the templates shown in Figures 2 and 3 with the detector thresholds given in (31). The SUMO detector also uses a tiling approach to calculating the background statistics rather than moving windows. The latter is probably done for computational efficiency but we believe this is not necessary. The moving window approach is more accurate and can be efficiently implemented as described by Wackerman. **We recommend that both the AKDEMO detector and the SUMO detector (with moving windows) be tried.** Although, we note that these detectors are very similar and it is expected that performance will likewise be similar. We also add that it should be easy to convert one to the other.

SAR system resolution and ship size need to be taken into account when choosing the detector window sizes. For the AKDEMO system, Wackerman *et al.* recommend that the target window be a similar size to smallest ship, the guard window a similar size to the largest ship, and the background window should be big enough to get accurate background statistics but small enough to be local. Similarly, SUMO's template should be adjusted to take account of resolution. Assuming equal range and azimuth resolution and random ship orientation the windows should be square. Of course, if this not the case then other window shapes should be considered. The actual window sizes used in the AKDEMO system were discussed in Section 5.3. Wackerman *et al.* note that ideally a dynamic target window, which adapts to the size and shape of the target to be detected, should be used but does not propose a technique for doing so. One possible alternative is to use several different target windows and pool the resulting detections. This probably explains the two different templates used by the SUMO system.

Setting design parameter t in (15) is the most important aspect of implementing a two parameter CFAR detector. This can be done theoretically to achieve a desired FAR, but as Wackerman *et al.* note, several different background models are

applicable, and it is probably better done empirically using trial and error. Commonly $t = 5.0$ is chosen (for single pixel target windows) but larger values such as $t = 6.0$ are also used. Wackerman *et al.* use larger values of t in higher resolution imagery. Note that when multi-pixel target windows are used, it is important to adjust t to $t/\sqrt{n}$ where n is the number of independent samples (which is not necessarily the same as the number of pixels) in the target window. In the case, of the SUMO detector there are 3 design parameters to adjust — the factors 3.0, 1.9 and 1.4 in (31). It is of interest to note here that $t = 5.0$ for a single pixel target window converts roughly to $t = 2.5$ for a 4 pixel target window and that this is close to the value 3.0 used in SUMO. In other words, the SUMO template is akin to using a 4 pixel target window. The JRC researchers have commented that the factors 3.0, 1.9 and 1.4 were found empirically and that the best setting for them varies from image to image. They have indicated their intention to make the adjustment process automatic but no details have appeared yet. Finally, we note that the setting of threshold design parameters will depend to some extent on the effectiveness of later false alarm discrimination stages. Better false alarm discrimination means a higher FAR can be tolerated at the prescreener stage which means that a lower threshold can be used and better probability of detection achieved.

Given the high profile of the OMW and the extensive testing it has under gone, a K-distribution based CFAR detector also warrants serious consideration. Further, the OMW grew out of the Ocean Features Workstation (OFW) which used a two parameter CFAR detector with $t = 5.0$. Rey *et al.*, [103], report that the FAR was too high in the OFW and better results were obtained with the K-distribution based CFAR detector. On the other hand, while modelling ocean backscatter with a K-distribution is well justified theoretically, it does not always seem to be appropriate in practice. It has been observed that the K-distribution is not always a good fit to SAR ocean data by Rey *et al.*, [102], Henschel *et al.*, [49], and Yeremy *et al.*, [143]. In circumstances where the K-distribution fails (as revealed by a negative order parameter) the OMW uses a chi-squared distribution (more commonly referred to as a gamma distribution) instead. **Consequently, a gamma distribution based CFAR detector where both the mean and the order parameter are estimated from the data should also be considered.** Details of the OMW detector were discussed in Section 5.4. We propose only one alteration. The OMW uses a tiling approach to calculating the background statistics for computational efficiency but again, **a moving window approach is likely to be more accurate.** A moving window K-distribution based CFAR detector was used by Robertson *et al.*, [105].

As we noted in Section 5.3, the detector used in the Eldhuset system, [32], is effectively the cell-averaging CFAR detector, (16). This detector is also used in the Alenia Areospazio ship detection system described by Ferrara *et al.* in [33] and with modifications in the system described by Cusano in [25]. The theoretical basis for the cell-averaging CFAR is gamma (or exponential) distributed background clutter. In this case $\mu_b = \sigma_b\sqrt{L}$ where μ_b is the background mean, σ_b is the background variance and L is the gamma order parameter (number of looks) and so the cell-averaging CFAR is the same as the two parameter CFAR but with a different value of the design parameter t to achieve the same FAR. We believe the two parameter CFAR is likely to be a better detector since it specifically takes account of the background variance and hence is applicable to other background distributions as well. As confirmation of this we note that Cusano *et al.*

modify the basic cell-averaging CFAR detector (to cope with high background variability) by allowing the design parameter in (16) to depend directly on the background variance. Similarly, Ferrara *et al.* comments that the CA-CFAR is unreliable in non-homogeneous areas. Consequently, we do not recommend the cell-averaging CFAR detector. Having said that we add that a two parameter CFAR can be converted to a cell-averaging CFAR easily and so both can be tried without extra effort if desired.

The Alenia Areospazio system described by Ferrara *et al.* in [33] also makes use of an order statistics CFAR detector. Recommending this detector is questionable since it is computationally expensive and Ferrara *et al.* have not provided convincing documentation of its benefits. On the other hand, we note that many researchers have reported detection problems in image areas with high variability and order statistics are generally recognised as being robust when applied to highly non-Gaussian data. Further, the Alenia Areospazio system only applies this detector to non-homogeneous areas of the image and so the computational expense is less than otherwise may be expected. As an aside, we note that an area is considered to be inhomogeneous if $\sigma_b > \mu/\sqrt{L}$. Our recommendation then is that **the order statistics CFAR detector should only be considered if it is found that false alarms in inhomogeneous areas is a significant problem.**

The G-GLRT detector is the only detector to use a target model (albeit a simplistic one). The tests in [117] indicate that this detector outperforms a cell averaging CFAR detector but they are both outperformed by a two parameter CFAR detector. However, the results also show that if an initial segmentation stage is included so that the detectors are applied to homogeneous regions only then the G-GLRT detector can outperform the two parameter CFAR detector. Given that the G-GLRT detector just uses target and background means and that it should therefore be relatively easy to adapt a two parameter CFAR detector into a G-GLRT detector, **we recommend that this detector be tested, especially if an initial segmentation stage is included.**

Next we mention, the Lockheed Martin Canada detector described by Gagnon *et al.* in [41]. While the performance of this detector has only been reported for land imagery, we believe it may be useful for ship detection purposes. Indeed, it is apparent from the report, [41], that that is one of the intended applications. This detector is especially attractive since it is the only one discussed in this report which is specifically aimed at high resolution imagery. Further, its use of the discrete wavelet transform incorporates not only multiresolution analysis but also "edge enhancment". We note too that, according to Gagnon *et al.*, this detector has a much lower FAR than the usual CFAR method of detection. Consequently, **we recommend that this detector be tried if high resolution imagery is of interest.** We add however that other detectors can sensibly be applied to high resolution imagery by down sampling the imagery first or otherwise reducing the resolution.

The detector proposed by Jiang *et al.*, [54], is of interest since it is CFAR based on a non-parametric estimation of the ocean background distribution. This work grew out of earlier work with a K-distribution based CFAR detector, [55], where it was found that the K-distribution was not always a good fit to the data. The non-parametric technique used is essentially a Parzen window method of approximating the pdf from histogram data. Jiang *et al.* note that the problem of estimating the correct "kernel width" has not been completely resolved and we anticipate difficulties here. Further, the method will be

difficult to implement and its performance has been poorly validated. Consequently, we only give a qualified recommendation of this detector.

The DSTO detectors reported on by Hunt *et al.*, [51], show promise. Their performance was shown to be substantially better than the two parameter CFAR detector. However, Hunt *et al.* note that the test data was limited and that it is difficult to draw conclusions from the results. Consequently, we recommend further testing of these detectors on more representative data, although we note that they are likely to be computationally expensive.

In earlier work on ship detection algorithms, Eldhuset, [30], used an adaptive Weiner filter to reduce speckle and then a Butterworth highpass filter to enhance the ship backscatter. The final detection was done with a simple threshold. In the later paper, [32], reports that this original detection scheme was too slow and hence was abandoned. It is presumed that the loss of accuracy (if any) in the new faster detector was not significant. Further, no-one has followed up on Eldhuset's original detector. For these reasons, we do not recommend Eldhuset's original scheme.

10.6 Improving detector performance

Often the basic detectors are modified to improve detection performance. We mention the more useful modifications here. Once a ship detection system has been implemented and its performance analysed, the relevant modifications can be considered.

Wackerman *et al.*, [135], note that excessive false alarms can be a problem in areas with a large number of noise floor samples. One of his solutions is to increase the threshold design parameter t in such areas. Wackerman *et al.* also use special processing in the buffer zone around the land mask in order to avoid excessive false alarms from land backscatter. This is done by placing constraints on the background mean μ_b and standard deviation σ_b. Rey *et al.*, [102], note that background statistics can be corrupted by the presence of bright targets and recommends that a method of excluding unusually bright pixels be considered. We note that this problem is exacerbated in the OMW by its tiling approach to calculating background statistics. One of the solutions Rey *et al.* offer for the OMW detector, is the use of a first pass with an empirical threshold to identify the bright pixels. This idea is also proposed by Casasent *et al.*, [19], for the standard two parameter CFAR detector. Casasent *et al.* also note that false alarms can occur in the sidelobes of very bright objects and proposes the use of a sidelobe filter to mask out the sidelobes. Several other modifications dealing with various detection problems are also offered by Casasent *et al.*. These are all aimed at reducing the false alarm rate but are somewhat *ad hoc*. Finally, we note that Cusano *et al.*, [25], identify regions of very high contrast and masks them out. Unfortunately, Cusano *et al.* provide no details of how contrast is measured.

10.7 Polarimetric detectors

The topic of ship detection using polarimetric data is relatively new and research is ongoing. Not many well validated results are available in the literature. Consequently, **making recommendations in this area is particularly difficult.** However, it is anticipated that the use of polarimetric data will improve ship detection performance (see

for instance the comments by Yeremy *et al.* in [142]) and algorithms are currently being developed. Therefore we see this as an area where future research effort would be well spent.

In several papers, see for instance [124] and [123], Touzi *et al.*, have investigated the use of polarisation entropy for ship detection. The polarisation entropy they refer to is similar to and was motivated by Cloude and Pottier's target decomposition method, [22], which extracts entropy and an α-angle from the polarimetric data. While Touzi *et al.* claim to have a new definition of polarisation entropy, they have not yet provided the details. (In a private communication, Touzi has indicated that he intends to publish the details later this year.) Their results show that using polarisation entropy is an effective way of enhancing the contrast between ships and the ocean. **Hence Touzi's polarisation entropy is of interest for ship detection and should be followed up when the details become available.** We also add here the comment by Yeremy *et al.* in [142], that preliminary studies at Satlantic have shown the Cloude and Pottier decomposition to be a good method for ship detection. The effectiveness of this decomposition for target detection on land has recently been demonstrated at DSTO by Williams and Blucher [141]. **Thus the Cloude and Pottier decomposition may be worth further investigation too.** Note that the calculation of polarisation entropy is really a preprocessing technique and it needs to be followed by a bright anomaly detector.

Sciotti *et al.*, [116] and [117], list three fundamentally different ways of handling multi-channel polarimetric data.

A) Apply single channel detectors to each polarimetric channel separately and then fuse the detection results;

B) Fuse the polarimetric channels into a single channel and then apply single channel detectors;

C) Apply multi-channel detection algorithms.

However, apart from the reports of Sciotti *et al.*, not much work has been undertaken in the research community to compare these various strategies for ship detection purposes. Further, the results in [116] are mostly from simulated data and the results in [117] do not include an assessment of approach A). We therefore believe that **this area needs to be revisited and more extensive testing undertaken.** It is evident from [117] that a good baseline to start with is strategy B) with the use of the 2P-CFAR on SPAN or PWF filtered imagery. The 2P-CFAR clearly outperformed both the CA-CFAR and the G-GLRT on such imagery. Consequently, in the absence of more complete information **we recommend approach B) with the use of a 2P-CFAR on SPAN or PWF filtered imagery with PWF filtering likely to be slightly better.**

The only multi-channel detection algorithm considered by Sciotti *et al.* in approach C) was the PG-GLRT. While the earlier paper, [116], was optimistic about its performance, the later paper, [117], shows its performance to be variable. In one of the two images used, PG-GLRT outperformed the 2P-CFAR on SPAN filtered imagery well, but in the other its performance was considerably worse. **So the PG-GLRT is of some interest and we recommend that it be investigated further.** However, we warn that it is likely to be

difficult to implement and computationally expensive (matrix determinants and inverses are needed).

The Cameron decomposition is aimed at determining the dominant elemental scatterer type in each image pixel based on physical scattering mechanisms. Research by Ringrose *et al.*, [104], and Yeremy *et al.*, [142], has shown that in general the elemental scatterers on ships differ to those in the ocean and so the Cameron decomposition should be useful for ship detection purposes. The results from both papers indicate that the method has good potential for ship detection purposes. The study by Yeremy *et al.*, being much more extensive is particularly convincing, and their conclusion is that the Cameron method is a very robust ship detection method. They add that when a threshold method is incorporated with the Cameron method, detection rates are higher than with either method alone. Unfortunately, both the studies are more investigative than developmental and fully automatic algorithms have not yet been designed. **This is clearly an area for future research.**

The SSCM (symmetric scattering characterisation method) of Touzi and Charbonneau, [124] and [120] , is an extension of the ideas behind the Cameron method. Like the work of Ringrose and Yeremy, Touzi and Charbonneau's work is investigative and suggests avenues for further exploration rather than polished algorithms. It should also be noted that their main interest is in ship identification and characterisation. Still, their approach involves several improvements to the basic Cameron method and hence should be applicable to ship detection too. If research effort is to be invested into the Cameron method then **the SSCM should be considered at the same time.**

10.8   Segmentation

In [70], Lombardo and Sciotti note that the CA-CFAR and related detectors struggle in ocean areas which are inhomogeneous. They note in particular that while taking the background variance into account and raising the detection threshold in high contrast areas reduces the FAR, it also reduces the probability of detection. They propose an initial segmentation step as a way of overcoming such problems. The idea is that segmentation can identify step edges and other such inhomogeneities and hence allows the CFAR detectors to calculate background statistics from the same homogeneous region as the pixel under test. Evidently this applies to any of the adaptive threshold algorithms already discussed. We add here Lombardo's use of segmentation is also effectively a way of adapting the CA-CFAR threshold to varying image contrast. Lombardo and Sciotti's experiments, [70], with single channel SAR imagery and the CA-CFAR detector show that segmentation improves detection performance substantially. In later work, [117], Sciotti *et al.* considered multi-channel polarimetric imagery. Again, it was found that prior segmentation led to significant improvements in performance.

Thus prior segmentation is an attractive proposal for ship detection. On the down side, some aspects of the two algorithms described for single channel imagery in [70] are *ad hoc* and several of the details are not reported. Further, different segmentation and detection algorithms are used in low and high resolution imagery. It seems likely then that details of the algorithms are dependent on the SAR imaging parameters and that adjustments would be needed in order to apply the algorithm to new imagery. Similarly, while

various approaches to segmentation of multi-channel polarimetric imagery are discussed in [117] the details are not really mentioned at all.

Thus, while Lombardo and Sciotti's work is conceptually appealing and is backed up with convincing performance results, we believe it may not be robust and that effort would be needed to fine tune it to new surveillance scenarios. We add that accurate segmentation is hard and computationally expensive. In other words, **we see Lombardo and Sciotti's contribution as more of an indication of potential rather than a completed algorithm and we view it as an area for future research which shows good promise.**

10.9 Interferometric SAR detectors

Along-track interferometric SAR (InSAR) can be used to detect moving targets on the ocean surface including ships and ship wakes. This is done by estimating velocities from the phase information in the InSAR interferogram. Since these velocity measurements do not depend on RCS, small ships which are not much brighter than the background ocean clutter can be detected. We note however, that along-track interferometric SAR (InSAR) involves fundamental changes to the radar setup — a second antenna must be added. Also, while the effectiveness of the technique was demonstrated by Campbell *et al.* [15], no automatic processing of the imagery was undertaken. **Thus, along-track interferometric SAR is an attractive proposition but it requires a specialised radar system and detection algorithms still need to be developed.**

A somewhat different technique of using SAR interferometry for ship detection has been described by Arnaud in [6]. Again, the idea is to use phase rather than intensity to aid in the detection of small ships with low RCS. However, Arnaud proposes to form the interferogram from two different looks processed from a single aperture SAR system. The idea here is that there is a time delay between the two looks and the ocean which is continuously changing shape will have strongly incoherent behaviour while ships will tend to have more coherent behaviour. Thus Arnaud detects potential ships by applying a simple threshold to the phase coherence image. **Given that this interferometric SAR method does not involve a specalised SAR system and should allow the detection of vessels with a low RCS it is of interest.** Ideas very similar to Arnaud's have been developed in work of Iehara *et al.*, [53], and Ouchi *et al.*, [93], — the second paper being an extension of the work in the first. However, while Arnaud, Iehara and Ouchi all demonstrate the feasibility of this approach, their tests are limited and **the practical utility of the method has not yet been demonstrated.**

10.10 Clustering and elementary discrimination

There is not a lot to say on this topic. **Most detection systems include a stage in which the individual detected pixels are clustered into extended detections and this is usually followed by some form of elementary discrimination.** Clustering is commonly done on the basis of pixel adjacency whereby contiguous detected pixels are assumed to arise from the same underlying object. Elementary discrimination is done on

the basis of simple measurements of the clustered detections. Clustering and elementary discrimination are recommended.

Both Wackerman *et al.*, [135], and Robertson *et al.*, [105], recommend a lower threshold be placed on the total intensity of the target backscatter to avoid false alarms in dark areas of the imagery. The OMW, see [102], rejects detections which are: too large (in terms of number of pixels); too small; too close to land; and too close to other targets. The JRC system also rejects detections which are too large. In order to avoid false alarms due to large objects with strong sidelobe patterns, the JRC system also clusters any detections which are too close to one another. Ship orientation can be calculated using standard techniques and a length and width extracted. Ship orientation is useful ancillary information for estimating the heading while length may used as a discriminating feature and an aspect ratio can be calculated from the length and width. It is noted by Askari and Zerr in [7] that the length-to-width ratio is a standard indicator of ship structure and for merchant ships this ratio is of the order of 6 to 7, whereas it is larger than 10 for military vessels. We also mention Gower and Skey's method, [43], of extracting target measurements by fitting a Gaussian profile. **The particular target measurements used for discrimination and their associated limits will be scenario dependent and should be left open for user specification** but care is needed since there will usually be a tradeoff of missed detections against false alarm reduction.

Some systems use morphological filters prior to clustering. Thus Lin *et al.*, [68], and Jiang *et al.*, [54], use an erosion filter to remove false alarms due to speckle. Similarly, Liu *et al.*, [69] propose both dilation and erosion. Dilation will connect nearby pixels and hence can be used for clustering. We also mention here the proposal in [107] to use two thresholds for clustering. The higher threshold is used as usual to detect likely ship pixels while the lower threshold is used to connect related detected pixels and hence cluster them. **Such filters should be easy to implement and test.** The only caveat on their use is that assessing their effect on performance will require extensive well ground-truthed data.

10.11 Discrimination

As we have noted before, not much progress has been made on discrimination algorithms for ship detection purposes. Both the OMW and the AKDEMO systems do not include any discrimination algorithms other than the elementary ones already discussed. However, **the SUMO and Eldhuset systems have discrimination algorithms which are worth trying.**

The JRC researchers, [114], report that many false alarms in the SUMO system are due to very local wind turbulence which produces very noisy backscatter conditions. We comment here that this problem may be exacerbated by the tiling approach used for calculating the detector thresholds. The JRC researchers note that one of the characteristics of these false alarms is a large number of local maxima above the primary detector threshold in the local area. Thus they propose counting the number of pixels with values above primary detector threshold and then rejecting the detection if this number is too high.

The Eldhuset system uses "ship-to-sea" condition test to discriminate false alarms. Eldhuset, [32], reports that this test is very necessary when swells, strong waves, eddies

and steep fronts are present and that it substantially reduces the number of false alarms. The test calculates the ratio of the integrated intensity of all ship pixels to the contrast of the local area. If this ratio is below a given threshold the detection is rejected. We also mention Eldhuset's homogeneity test. The idea of the homogeneity test is that the ocean area in front of a ship must be relatively homogeneous. While Eldhuset only applies this test when a potential wake has been detected, we note that it also makes sense in the absence of a wake as long as the ship heading can be determined. The details of this test are given in Section 8.

In [142] and [143], Yeremy *et al.* indicate that **polarimetric data has good potential for false alarm discrimination.** They refer in particular to the Cameron decomposition. The Cameron decomposition, see Section 6.6, classifies polarimetric pixels in terms of elemental scatterer types and hence provides information about scattering mechanisms and the underlying physical structures. This makes the Cameron decomposition suitable for discrimination and classification purposes. A good example is given by the finding of Yeremy *et al.* that the scatterer types (as revealed by the Cameron decomposition) associated with ocean buoys were very different to those of small ships. However, these findings have not yet been developed in to algorithms. Yeremy *et al.* note that further studies are required. We also mention here related work by Touzi and Charbonneau on the SSCM, see Section 6.7. But again, practical algorithms have not yet been developed.

Other discrimination techniques which have been used are the Kohonen neural network mentioned by Robertson *et al.* in [105] and the fuzzy classifier used by Benelli and Argenti in [8] and [5]. But these are not within DSTO's current technology base and they require extensive training data. Consequently, they are not recommended at this stage.

In connection with the need for training data we mention the work being undertaken at the Universitat Politecnica de Catalunya (UPC) by Mallorqui and others to generate simulated training data, see for example [80]. But this work is still in the developmental stage.

Wake detection is an obvious enhancement to a ship detection system. The presence of a wake confirms the presence of a ship and further, it can be used to estimate the ship heading and speed. However, as explained before, **wake detection cannot be used to confidently reject false alarms since wakes are often poorly imaged (especially when the imaging parameters are set to maximise sea-ship contrast) and stationary and slow moving ships have no wake anyway**. Thus, the best that can be done is to use the presence of a wake to confirm an otherwise weak detection. This is the outlook of Jiang *et al.*, [55], where the use of two detection thresholds is proposed. Detections above the higher threshold are accepted as is while detections above the lower threshold require a corroborating wake to be accepted.

10.12 Human supervision

Including a human in the loop to provide a final false alarm discrimination stage is important to several systems. In fact, Henschel *et al.*, [49], indicate that this is part of the OMW design philosophy. To quote them: *The false alarm rate is still an open question, however, our operational experience has shown that the person-machine mix, with*

very fast algorithms means that ships with specific type and classification characteristics can be supplied in a very short period of time. This means that in an operational scenario, the automated algorithm's false alarm rate can be ignored. The way this operates in practice is that the detections are ordered according to size (or some other appropriate criteria). The top 20 or so detections are then displayed to the human supervisor and a final assessment is made. Thus a GUI should be provided and relevant information displayed. As an example, the OMW displays the original SAR image with detections overlayed together with an enlargement of the particular detection being reviewed. Target location, width, length, speed and status are all listed. The status field is used by the user to classify the detections. Qinetiq's MaST system also displays a mask of the clustered pixels in the detection. Extra analysis such as wake detection can be made optional.

References

1. W. Alpers. Imaging ocean surface waves by synthetic aperture radar: a review. In T.D. Allan, editor, *Satellite Microwave Remote Sensing*, chapter 6. Ellis Horwood Ltd., West Sussex, England, 1983.

2. C. van Antwerpen and E. Kruzins. Commercial remote sensing satellites and their potential to perform maritime surveillance: results from the Satellite Shipping Study. Research Report DSTO-RR-0180, Defence Science & Technology Organisation, Australia, October 2000. 128 pages.

3. C. van Antwerpen and E. Kruzins. Spaced based synthetic aperture radar tool - a code and model for estimating the maritime detection capability of generic spaceborne synthetic aperture radars. Technical Report DSTO-TR-1050, Defence Science & Technology Organisation, Australia, September 2000. 32 pages.

4. E. Aresu and G. Schwartz. Impact of ScanSAR images' radiometric calibration on vessels and identification. In *IEEE 2003 International Geoscience and Remote Sensing Symposium (IGARSS'03)*, 2003.

5. F. Argenti, G. Benelli, A. Garzelli, and A. Mecocci. Automatic ship detection in SAR images. In *IEE International Conference Radar 92*, pages 465–8, 1992.

6. A. Arnaud. Ship detection by SAR interferometry. In *IEEE 1999 International Geoscience and Remote Sensing Symposium (IGARSS'99)*, volume 5, pages 2616–18, 1999.

7. F. Askari and B. Zerr. Automatic approach to ship detection in spaceborne synthetic aperture radar imagery: An assessment of ship detection capability using RADARSAT. Technical Report SACLANTCEN-SR-338, SACLANT Undersea Research Centre, La Spezia (Italy), December 2000. 36 pages.

8. G. Benelli, A. Garzelli, and A. Mecocci. Complete processing system that uses fuzzy logic for ship detection in SAR images. *IEE Proceedings: Radar, Sonar & Navigation*, 141(4):181–6, 1994.

9. D. Blacknell. Comparison of parameter estimators for K-distribution. *Radar, Sonar and Navigation, IEE Proceedings*, 141(1):45–52, 1994.

10. M. Brizi, P. Lombardo, and D. Pastina. Exploiting the shadow information to increase the target detection performance in SAR images. In *International Conference on Radar Systems, RADAR'99*.

11. R.J. Brown. Ship detection. *Backscatter*, (Fall):15–15, 2000.

12. W.L. Cameron and L.K. Leung. Feature motivated polarization scattering matrix decomposition. In *Proceedings of the IEEE 1990 International Radar Conference (Radar'90)*, pages 549–57, 1990.

13. W.L. Cameron, N.N. Youssef, and L.K. Leung. Simulated polarimetric signatures of primitive geometrical shapes. *IEEE Transactions on Geoscience and Remote Sensing*, 34:793–803, 1996.

14. J.W. Campbell, P.W. Vachon, and F.W. Dobson. The extraction of ocean surface information from SAR imagery. In *International Symposium, Geomatics in the Era of RADARSAT (GER'97), 15 pages*, 1997. Ottawa, Canada, May 25-30.

15. J.W.M. Campbell, A. L. Gray, K.E. Mattar, J.H. Clarke, and M.W.A. van der Kooij. Ocean surface feature detection with the CCRS along-track InSAR. *Canadian Journal of Remote Sensing*, 23(1):24–37, 1997.

16. J.W.M. Campbell and P.W. Vachon. Optimisation of the Ocean Features Workstation for ship detection. Technical report, CCRS/AD/ADS, 11 July 1997. 15 pages.

17. R.E. Carande. Estimating ocean coherence time using dual-baseline interferometric synthetic aperture radar. *IEEE Transactions on Geoscience and Remote Sensing*, 32(4):846–854, 1994.

18. W.G. Carrara, R.S. Goodman, and R.M. Majewski. *Spotlight Synthetic Aperture Radar: Signal Processing Algorithms.* Artech House, Norwood, Massachusetts, USA, 1995.

19. D. Casaent, Wei Su, Deepak Turaga, Nobuhiko Narusawa, and Satoshi Ashizawa. SAR ship detection using new conditional contrast box filter. In *SPIE Conference on Algorithms for Synthetic Aperture Radar Imagery VI*, volume 3721, pages 274–84, April 1999.

20. P.B. Chapple, D.C. Bertilone, R.S. Caprari, and G.N. Newsam. Stochastic model-based processing for detection of small targets in non-Gaussian natural imagery. *Image Processing, IEEE Transactions on*, 10(4):554–564, 2001.

21. Jinsong Chong, Minhui Zhu, and Ge Dong. Ship target segmentation of high-resolution SAR images. In *4th European Conference on Synthetic Aperture Radar (EUSAR2002)*, 2002.

22. S.R. Cloude and E. Pottier. A review of target decomposition theorems in radar polarimetry. *IEEE Transactions on Geoscience and Remote Sensing*, 34:498–518, 1996.

23. S.R. Cloude and E. Pottier. An entropy based classification scheme for land applications of polarimetric SAR. *IEEE Transactions on Geoscience and Remote Sensing*, 35:68–78, 1997.

24. R. Cook, I. McConnell, D. Stewart, and C. Oliver. Segmentation and simulated annealing. In *Proceedings of SPIE - Microwave Sensing and Synthetic Aperture Radar*, volume 2958, pages 30–7, 1996.

25. M. Cusano, J. Lichtenegger, P. Lombardo, A. Petrocchi, and D. Zanovello. A real time operational scheme for ship traffic monitoring using quick look ERS SAR images. In *IEEE 2000 International Geoscience and Remote Sensing Symposium (IGARSS'00)*, volume 7, pages 2918–20, 2000.

26. M. Dinguirard, T. Wahl, J.M. McDonald, G. Carrin, D. Levy, U. Erich, K.H. Zeller, and X. Briottet. Some results from AGARD study AAS-42: Potential use of commercial satellites systems for military applications. In *AGARD-CP-582, AGARD SPP*

Symposium on "Remote Sensing: A valuable source of information", 1996. Toulouse, France, 22–25 April.

27. M. Dowd, P.W. Vachon, F.W. Dobson, and R.B. Olsen. Ocean wave extraction from Radarsat synthetic aperture radar inter-look image cross-spectra. *IEEE Transactions on Geoscience and Remote Sensing*, 39(1):21–37, 2001.

28. Kruzins E. The effectiveness of commercial satellites for surveillance, search and rescue. *Journal of Electrical & Electronic Engineering, Australia*, 20(3):147–170, 2000.

29. C. Elachi. *Spaceborne Radar Remote Sensing: Applications and Techniques*. IEEE Press, New York, USA, 1988.

30. K. Eldhuset. Automatic ship and ship wake detection in spaceborne SAR images from coastal regions. In *IEEE 1988 International Geoscience and Remote Sensing Symposium (IGARSS'88)*, pages 1529–1533, 1988.

31. K. Eldhuset. Principles and performance of an automated ship detection system for SAR images. In *IEEE 1989 International Geoscience and Remote Sensing Symposium (IGARSS'89)*, pages 358–361, 1989.

32. K. Eldhuset. An automatic ship and ship wake detection system for spaceborne SAR images in coastal regions. *IEEE Transactions on Geoscience & Remote Sensing*, 34(4):1010–19, 1996.

33. M.N. Ferrara, A. Gallon, and A. Torre. Improvement in automatic detection and recognition of moving targets in Alenia Aerospazio activity. In *Proceedings of SPIE - EUROPTO Conference on Image and Signal Processing for Remote Sensing*, volume 3500, pages 96–103, September 1998.

34. M.N. Ferrara and A. Torre. Automatic moving targets detection using a rule-based system: comparison between different study cases. In *IEEE 1998 International Geoscience and Remote Sensing Symposium (IGARSS'98)*, volume 3, pages 1593–1595, 1988.

35. M.N. Ferrara and A. Torre. Alenia Aerospazio activity on SAR data analysis and information extraction. In *Proceedings of SPIE*, volume 3217, pages 67–75, 1997.

36. M.F. Fingas. Detection of ships from airborne platforms: a summary of workshop presentations. *Backscatter*, (Fall):32–35, 2000.

37. M.F. Fingas and C.E. Brown. Review of ship detection from airborne platforms. *Canadian Journal of Remote Sensing*, 27(4):379–385, 2001.

38. K.S. Friedman, C. Wackerman, F. Funk, W.G. Pichel, P. Clemente-Colon, and Xiaofeng Li. Validation of a CFAR vessel detection algorithm using known vessel locations. In *IEEE 2001 International Geoscience and Remote Sensing Symposium (IGARSS'01)*, volume 4, pages 1804–6, 2001.

39. K.S. Friedman, C. Wackerman, F. Funk, K. Rowell, W.G. Pichel, P. Clemente-Colon, and Xiaofeng Li. Validation of an automatic vessel detection algorithm using SAR data and known vessel fleet distributions. In *IEEE 2000 International Geoscience and Remote Sensing Symposium (IGARSS'00)*, volume 5, pages 2071–3, 2000.

40. V.S. Frost, J.A. Stiles, K.S. Shanmugan, and J.C. Holtzman. A model for radar images and its application to adaptive digital filtering of multiplicative noise. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 4:157–66, 1982.

41. L. Gagnon, H. Oppenheim, and P. Valin. R&D activities in airborne SAR image processing/analysis at Lockheed Martin Canada. In *Proceeding of SPIE*, volume 3491, pages 998–1003, 1998.

42. G. B. Goldstein. False-alarm regulation in log-normal and Weibull clutter. *IEEE Transactions on Aerospace and Electronic Systems*, 9(1):84–92, 1973.

43. J. Gower and S. Skey. Evaluation of RADARSAT ScanSAR for observing wind, slicks and fish-boats. *Canadian Journal of Remote Sensing*, 26(6):484–493, 2000.

44. J. Gower and S. Skey. Wind, slick and fishing boat observations with Radarsat ScanSAR. *Johns Hopkins APL Technical Digest*, 21(1):68–74, 2000.

45. R.K. Hawkins, K.P. Murnaghan, T. Tennant, M. Yeremy, and M. Rey. Ship detection using airborne polarimetric SAR. In *CEOS SAR Workshop 2001*, 2–5 April 2001. Tokyo, Japan.

46. R.K. Hawkins, R. Touzi, P.W. Vachon, J.R. Gibson, and M. van der Kooij. Data quality requirements for future SAR. In *Proceedings of the 22nd Canadian Symposium on Remote Sensing*, pages 119–124, 2000.

47. M.D. Henschel, P.A. Hoyt, J.H. Stockhausen, P.W. Vachon, M.T. Rey, J.W.M. Campbell, and H.R. Edel. Vessel detection with wide area remote sensing. *Sea Technology*, 39(9):63–68, 1998.

48. M.D. Henschel, R.B. Olsen, P. Hoyt, and P.W. Vachon. The Ocean Monitoring Workstation: experience gained with RADARSAT. In *Proceedings, Geomatics in the Era of RADARSAT'97*, 25–30 May 1997. Ottawa, Canada.

49. M.D. Henschel, M.T. Rey, J.W.M. Campbell, and D. Petrovic. Comparison of probability statistics for automated ship detection in SAR imagery. In *Proceedings of SPIE*, volume 3491, pages 986–91, 1998.

50. J.K. Hsiao. On the optimisation of MTI clutter rejection. *IEEE Transactions on Aerospace and Electronic Systems*, 10(5):622–629, 1974.

51. E. Hunt, S. Angeli, G.N. Newsam, and E. Kruzins. Spase-based synthetic aperture radar surveillance: Detection of small targets in a maritime environment. Technical Report DSTO-TR-0980, Defence Science & Technology Organisation, Australia, February 2001. 56 pages.

52. J.R. Huynen. Phenomenological theory of radar targets. In P.L. Uslenghi, editor, *Electromagnetic scattering*. New York: Academic, 1978.

53. M. Iehara, K. Ouchi, I. Takami, K. Morimura, and S. Kumano. Detection of ships using cross-correlation of split-look SAR images. In *IEEE 2001 International Geoscience and Remote Sensing Symposium (IGARSS'01)*, volume 4, pages 1807–9, 2001.

54. Qingshan Jiang, E. Aitnouri, S. Wang, and D. Ziou. Automatic detection for ship target in SAR imagery using PNN-model. *Canadian Journal of Remote Sensing*, 26(4):297–305, 2000.

55. Qingshan Jiang, Shengrui Wang, Djemel Ziou, Ali El Zaart, M.T. Rey, G.B. Benie, and M. Henschel. Ship detection in RADARSAT SAR imagery. In *IEEE 1998 International Conference on Systems, Man, and Cybernetics (SMC'98)*, volume 5, pages 4562–6, 1998.

56. J.A. Johannessen. Coastal observing systems: The role of synthetic aperture radar. *Johns Hopkins APL Technical Digest*, 21(1):41–48, 2000.

57. Joint Research Centre (JRC) of the European Community — IMPAST and DECLIMS web site, http://intelligence.jrc.cec.eu.int/marine/fish/index.htm.

58. N. Kourti and I. Shepard. Marine surveillance case study — integrating active and passive satellite-based technologies to improve European fisheries monitoring and control. http://www.rsi.ca/rsic/marine/mar_euro_jrc1.asp.

59. N. Kourti, I. Shepard, D. Brock, S. Moessel, P. Griffith, and D. Martin. Ship detection for fisheries monitoring: a summary of workshop presentations. *Backscatter*, (Fall):16–20, 2000.

60. N. Kourti, I. Shepard, and J. Verborgh. Fishing boat detection by using SAR imagery. In *Ship Detection in Coastal Waters Three Day Workshop*, 2000.

61. N. Kourti, I. Shepherd, G. Schwartz, and P. Pavlakis. Integrating spaceborne SAR imagery into operational systems for fisheries monitoring. *Canadian Journal of Remote Sensing*, 27(4):291–305, 2001.

62. E. Kruzins, Y. Dong, and B.C. Forster. Detection of vessels using SAR for support to search and rescue. In *9th Australasian Remote Sensing and Photogrammetry Conference*, volume 1, page 7304, 1998.

63. V. Larson, L. Novak, and C. Stewart. Joint spatial-polarimetric whitening filter to improve SAR target detection performance for spatially distributed targets. In *Proceedings of SPIE - Algorithms for synthetic aperture radar imagery*, volume 2230, pages 285–301, 1994.

64. Jong-Sen Lee and I. Jurkevich. Coastline detection and tracing in SAR images. *IEEE Transactions on Geoscience and Remote Sensing*, 28(4):662–668, 1990.

65. F. Lesage and L. Gagnon. Experimenting level set-based snakes for contour segmentation in radar imagery. In *Proceedings of SPIE - Visual Information Processing IX*, volume 4041, pages 154–62, 2000.

66. H. Leung, N. Dubash, and N. Xie. Detection of small objects in clutter using a GA-RBF neural network. *IEEE Transactions on Aerospace & Electronic Systems*, 38(1):98–118, 2002.

67. I-I. Lin and V. Khoo. Computer-based algorithm for ship detection from ERS-XSAR imagery. In *Proceedings of the 3rd ERS Scientific Symposium, Florence, Italy, 17-21 March*, 1997.

68. I-I. Lin, Leong Keong Kwoh, Yuan-Chung Lin, and V. Khoo. Ship and ship wake detection in the ERS SAR imagery using computer-based algorithm. In *IEEE 1997 International Geoscience and Remote Sensing Symposium (IGARSS'97)*, pages 151–3, 1997.

69. Yingjian Liu, Mingqiang Fang, Qian Feng, and Laibu Wang. An automatic ship detection system using ERS SAR images. In *IEEE 2003 International Geoscience and Remote Sensing Symposium (IGARSS'03)*, 2003.

70. P. Lombardo and M. Sciotii. Segmentation-based technique for ship detection in SAR images. *IEE Proceedings: Radar, Sonar & Navigation*, 148(3):147–59, 2001.

71. P. Lombardo, M. Sciotti, and L.M. Kaplan. SAR prescreening using both target and shadow information. In *IEEE National Radar Conference - Proceedings 2001*, pages 147–152, 2001.

72. M.W. Long. *Radar Reflectivity of land and sea*. Artech House, London, UK, 3rd edition, 2001.

73. A. Lopes, J. Bruniquel, F. Sery, J-C. Souyris, and F. Adragna. Optimal target detection using one channel SAR complex imagery: application to ship detection. In *IEEE 1998 International Geoscience and Remote Sensing Symposium (IGARSS'98)*, volume 2, pages 813–15, 1998.

74. A. Lopes, E. Nezry, R. Touzi, and H. Laur. Structure detection and statistical adaptive speckle filtering in SAR images. *International Journal of Remote Sensing*, 14:1735–1758, 1993.

75. Liming Lu, Runsheng Wang, and Wugao Li. A target detection method in Range-Doppler Domain from SAR echo data. In *16th International Conference on Pattern Recognition (ICPR'02)*, volume 1, pages 91–94, 2002.

76. A.P. Luscombe and L. Lightstone. Adapting operations of the Radarsat SAR to enhance ship monitoring. In *IEEE 1997 International Geoscience and Remote Sensing Symposium (IGARSS'97)*, volume 1, pages 154–6, 1997.

77. J.J. Mallorqui, J.M. Rius, and M. Bara. Satellite SAR simulator for fishing vessels signature studies. In *4th European Conference on Synthetic Aperture Radar (EU-SAR2002)*, 2002.

78. J.J. Mallorqui, J.M. Rius, and M. Bara. Simulation of polarimetric SAR vessel signatures for satellite fisheries monitoring. In *IEEE 2002 International Geoscience and Remote Sensing Symposium (IGARSS '02)*, pages 2711–13, 2002.

79. G. Margarit, P. Blanco, J. Sanz, and J.J. Mallorqui. Orbital SAR simulator of fishing vessel polarimetric signatures based on high frequency electromagnetic calculations. In *IEEE 2003 International Geoscience and Remote Sensing Symposium (IGARSS '03)*, 2003.

80. G. Margarit, P. Blanco, J. Sanz, J.J. Mallorqui, and J.M. Rius. Simulation of fishing vessels polarimetric signatures as first step to vessel classification. In *SPIE Conference on Image Analysis, Modeling and Techniques VI*, volume 5236, pages 154–62, 2004.

81. I. McConnell and C. Oliver. Segmentation-based target detection in SAR. In *Proceedings of SPIE - SAR image analysis, modeling and techniques II*, volume 3869, pages 45–54, 1999.

82. D.R. Montgomery. International fisheries enforcement management using wide swath SAR. *Johns Hopkins APL Technical Digest*, 21(1):141–147, 2000.

83. M.-C. Mouchot and R. Garello. SAR for oceanography. In F. M. Henderson and A.J. Lewis, editors, *Manual of Remote Sensing, Volume 2: Principles and Applications of Imaging Radar*, chapter 12. John Wiley & Sons, Inc., New York, USA, 3rd edition, 1998.

84. L. M. Novak, G. J. Owirka, and C. M. Netishen. Performance of a high-resolution polarimetric SAR automatic target recognition system. *Lincoln Laboratory Journal*, 6(1):11–24, 1993.

85. L.M. Novak and M.C. Burl. Optimal speckle reduction in polarimetric SAR imagery. *IEEE Transactions on Aerospace and Electronic Systems*, 26(2):293–305, 1990.

86. C. Oliver and S. Quegan. *Understanding Synthetic Aperture Radar Imagery*. Artech House, Norwood, Massachusetts, USA, 1998.

87. C.J. Oliver. Optimum texture estimators for SAR clutter. *Journal of Physics D-Applied Physics*, 26(11):1824–35, 1993.

88. C.J. Oliver and P. Lombardo. Simultaneous mean and texture edge detection in SAR clutter. *IEE Proceedings: Radar, Sonar & Navigation*, 143:391–9, 1996.

89. R.B. Olsen and T. Wahl. The role of wide swath SAR in high-latitude coastal management. *Johns Hopkins APL Technical Digest*, 21(1):136–140, 2000.

90. R.B. Olsen and T. Wahl. The ship detection capability of ENVISAT's ASAR. In *IEEE 2003 International Geoscience and Remote Sensing Symposium (IGARSS'03)*, 2003. Toulouse, France, July 21-25.

91. H. Osman and S.D. Blostein. Probabilistic winner-take-all segmentation of images with application to ship detection. *IEEE Transactions on Systems, Man, & Cybernetics, Part B: Cybernetics*, 30(3):485–90, 2000.

92. K. Ouchi, M. Iehara, K. Morimura, S. Kumano, and I. Takami. Nonuniform azimuth image shift observed in the Radarsat images of ships in motion. *Geoscience & Remote Sensing, IEEE Transactions on*, 40(10):2188–2195, 2002.

93. K. Ouchi and H. Yaguchi. Simulation on the extraction of ships' images embedded in speckle using cross-correlation of multilook SAR images and applications to Radarsat data. In *IEEE 2002 International Geoscience and Remote Sensing Symposium (IGARSS'02)*, volume 4, pages 2498 -500, 2002.

94. W.G. Pichel and P. Clemente-Colon. NOAA CoastWatch SAR applications and demonstration. *Johns Hopkins APL Technical Digest*, 21(1):49–57, 2000.

95. B. Plache and M.D. Henschel. Modification of the OMW to include polarimetric ship detection. Technical Report SAT-DN-0075, Satlantic Incorporated, Halifax, NS, 2001.

96. QinetiQ's ship detection system — Maritime Surveillance Tool (MaST) web site, http://www.qinetiq.com/casestudies/2002/case_study54.html.

97. R.K. Raney. SAR systems. In M. Ikeda and F.W. Dobson, editors, *Oceanographic Applications of Remote Sensing*, chapter Appendix E. CRC Press, Inc., New York, USA, 1995.

98. R.K. Raney, A.P. Luscombe, E.J. Langham, and S. Ahmed. RADARSAT. *Proceedings of the IEEE*, 79(6):839–849, 1991.

99. R. Ravid and N. Levanon. Maximum-likelihood CFAR for Weibull background. *Radar and Signal Processing, IEE Proceedings F*, 139(3):256–64, 1992.

100. N.J. Redding. Design of the Analysts' Detection Support System for broad area aerial surveillance. Technical Report DSTO-TR-0746, Defence Science & Technology Organisation, Australia, December 1998. 62 pages.

101. N.J. Redding, D.I. Kettler, G. Blucher, and P.G. Perry. The Analysts' Detection Support System for deploying a network of target detection and recognition algorithms in SAR exploitation. In *Proceedings of the 2003 IEEE Radar Conference (RADAR'03)*, 2003. Adelaide, Australia, September 3–5.

102. M.T. Rey, J. Campbell, and D. Petrovic. A comparison of ocean clutter distribution estimators for CFAR-based ship detection in RADARSAT imagery. Report No.1340, Defence Research Establishment Ottawa, Canada, December 1998. 31 pages.

103. M.T. Rey, A. Drosopoulous, and D. Petrovic. A search procedure for ships in RADARSAT imagery. Report No.1305, Defence Research Establishment Ottawa, Canada, December 1996. 35 pages.

104. R. Ringrose and N. Harris. Ship detection using polarimetric SAR data. In *CEOS SAR Workshop 2001*, 26–29 October 1999. Toulouse, France.

105. N. Robertson, P. Bird, and C. Brownsword. Ship surveillance using RADARSAT ScanSAR images. In *Alliance for Marine Remote Sensing (AMRS) Workshop on Ship Detection in Coastal Waters*, 2000.

106. I.S. Robinson. *Satellite Oceanography: an introduction for oceanographers and remote sensing scientists*. Ellis Horwood Ltd., Chichester, UK, 1985.

107. A. J. Rye, F. G. Sawyer, and R. Sothinathan. A workstation for the fast detection of ships. In *IEEE 1990 International Geoscience and Remote Sensing Symposium (IGARSS'90)*, volume 3, pages 2263–2266, 1990.

108. J.J. van der Sanden, P. Budkewitsch, D. Flett, A.L. Gray, R.K. Hawkins, Landry R., T.I. Lukowski, H. McNairn, T.J. Pultz, V. Singhroy, J. Sokol, T. Toutin, R. Touzi, and P.W. Vachon. RADARSAT-2; are its technical capabilities expected to provide potential for remote sensing applications? In *IEEE 2002 International Geoscience and Remote Sensing Symposium, 2002. (IGARSS '02)*, pages 1468–1470, 24–28 June 2002.

109. J.J. van der Sanden, P. Budkewitsch, D. Flett, R. Landry, T.I. Lukowski, H. McNairn, T.J. Pultz, V. Singhroy, J. Sokol, R. Touzi, and P.W. Vachon. Applications potential of planned C-band SAR satellites: leading to RADARSAT-2. In *IEEE 2001 International Geoscience and Remote Sensing Symposium, 2001. (IGARSS '01)*, pages 488–492, 9–13 July 2001.

110. J.J. van der Sanden and S.G. Ross. Applications potential of RADARSAT-2: A preview. Report, Canada Centre for Remote Sensing, 2001. 117 pages, unpublished but available at http://www.ccrs.nrcan.gc.ca/ccrs/rd/rd_e.html.

111. Satlantic Inc, Canada. *The Ocean Monitoring Workstation PC, Version 2001.1, User Manual Revision A*, March 2001. http://www.satlantic.com/products/by-name/omw/OMW_Manual.pdf.

112. K. Schulz, U. Soergel, and U. Thoennessen. Segmentation of moving objects in SAR-MTI data. In *Proceedings of SPIE - Algorithms for Synthetic Aperture Radar Imagery VIII*, volume 4382, pages 174–81, 2001.

113. K. Schulz, U. Soergel, U. Thoennessen, and U. Stilla. Elimination of across-track phase components in airborne along-track interferometry data to improve object velocity measurements. In *IEEE 2003 International Geoscience and Remote Sensing Symposium (IGARSS'03)*, 2003.

114. G. Schwartz, M. Alvarez, A. Varfis, and N. Kourti. Elimination of false positives in vessels detection and identification by remote sensing. In *IEEE 2002 International Geoscience and Remote Sensing Symposium (IGARSS'02)*, volume 1, pages 116–18, 2002.

115. M. Sciotti and P. Lombardo. Ship detection in SAR images: a segmentation-based approach. In *Proceedings of the 2001 IEEE Radar Conference (RADAR'01)*, pages 81–6, 2001.

116. M. Sciotti, D. Pastina, and P. Lombardo. Polarimetric detectors of extended targets for ship detection in SAR images. In *IEEE 2001 International Geoscience and Remote Sensing Symposium (IGARSS'01)*, volume 7, pages 3132–4, 2001.

117. M. Sciotti, D. Pastina, and P. Lombardo. Exploiting the polarimetric information for the detection of ship targets in non-homogeneous SAR images. In *IEEE 2002 International Geoscience and Remote Sensing Symposium (IGARSS'02)*, volume 3, pages 1911–13, 2002.

118. R. Touzi. On the use of polarimetric SAR data for ship detection. In *IEEE 1999 International Geoscience and Remote Sensing Symposium (IGARSS'99)*, volume 2, pages 812–14, 1999.

119. R. Touzi. Calibrated polarimetric SAR data for ship detection. In *IEEE 2000 International Geoscience and Remote Sensing Symposium (IGARSS'00)*, volume 1, pages 144–6, 2000.

120. R. Touzi and F. Charbonneau. Characterization of target symmetric scattering using polarimetric SARs. *IEEE Transactions on Geoscience and Remote Sensing*, 40:2507–2516, 2002.

121. R. Touzi and F. Charbonneau. The SSCM: an adaptation of Cameron's target decomposition to actual calibration SAR requirements. In *CEOS SAR Calibration Workshop Proceedings*, 2002. London, UK, September 23–26.

122. R. Touzi and F. Charbonneau. The SSCM for ship characterisation using polarimetric SAR. In *IEEE 2003 International Geoscience and Remote Sensing Symposium (IGARSS'03)*, 2003. Toulouse, France, July 21–25.

123. R. Touzi, F. Charbonneau, R.K. Hawkins, K. Murnaghan, and Kavoun X. Ship-sea contrast optimisation when using polarimetric SARs. In *IEEE 2001 International Geoscience and Remote Sensing Symposium (IGARSS'01)*, volume 1, pages 426–8, 2001.

124. R. Touzi, F. Charbonneau, R.K. Hawkins, and P.W. Vachon. Ship detection and characterization using polarimetric SAR. *Canadian Journal of Remote Sensing*, 2003. submitted.

125. F.T. Ulaby, R.K. Moore, and A.K. Fung. *Microwave Remote Sensing: Active and Passive, Volume II: Radar Remote Sensing and Surface Scattering and Emission Theory.* Addison-Wesley, Reading, Massachusetts, USA, 1982.

126. P.W. Vachon, P. Adlakha, H. Edel, M. Henschel, B. Ramsay, D. Flett, M. Rey, G. Staples, and S. Thomas. Canadian progress toward marine and coastal applications of synthetic aperture radar. *Johns Hopkins APL Technical Digest*, 21(1):33–40, 2000.

127. P.W. Vachon, J.W.M. Campbell, C.A. Bjerkelund, F.W. Dobson, and M.T. Rey. Ship detection by the RADARSAT SAR: Validation of detection model predictions. *Canadian Journal of Remote Sensing*, 23(1):48–59, 1997.

128. P.W. Vachon, J.W.M. Campbell, C.A. Bjerkelund, F.W. Dobson, and M.T. Rey. Detecting surface vessels by RADARSAT. In R.A. Brown, editor, *Remote sensing of the Pacific Ocean by Satellites*, pages 145–164. Earth Ocean & Space Pty. Ltd., 1998.

129. P.W. Vachon and F.W. Dobson. Wind retrieval from RADARSAT SAR images: Selection of a suitable C-band HH polarization wind retrieval model. *Canadian Journal of Remote Sensing, ADRO Final Symposium special issue*, 26(4):306–313, 2000.

130. P.W. Vachon and H.E. Krogstad. Synthetic aperture radar imagery of ocean waves. In M. Ikeda and F.W. Dobson, editors, *Oceanographic Applications of Remote Sensing*, chapter 15. CRC Press, Inc., New York, USA, 1995.

131. P.W. Vachon and R.B. Olsen. Radarsat SAR mode selection for marine applications. *Backscatter*, 6(3):3–4 & 18, 1995.

132. P.W. Vachon and R.B. Olsen. Ship detection with satellite-based sensors. *Backscatter*, (Fall):23–26, 2000.

133. P.W. Vachon, S.J. Thomas, J. Cranton, C. Bjerkelund, F.W. Dobson, and R.B. Olsen. Monitoring the coastal zone with the RADARSAT satellite. In *Oceanology International 98*, 10–13 March 1998.

134. P.W. Vachon, S.J. Thomas, J. Cranton, H.R. Edel, and M.D. Henschel. Validation of ship detection by the RADARSAT synthetic aperture radar and the Ocean Monitoring Workstation. *Canadian Journal of Remote Sensing*, 26(3):200–212, 2000.

135. C.C. Wackerman, K.S. Friedman, W.G. Pichel, P. Clemente-Colon, and X. Li. Automatic detection of ships in RADARSAT-1 SAR imagery. *Canadian Journal of Remote Sensing*, 27(5):568–577, 2001.

136. M.J. Wagner. From ships to shores to satellite images:. *Sea Technology*, 44(3):38–41, 2003.

137. T. Wahl, K. Eldhuset, and K. Aksnes. SAR detection of ships and ship wakes. In *Proceedings of the SAR Applications Workshop, Frascati, Italy. (ESA SP-264)*, pages 61–65, 16–18 September 1986.

138. T. Wahl, K. Eldhuset, and K. Skøelv. Ship traffic monitoring using the ERS-1 SAR. In *Proceedings First ERS-1 Symposium - Space at the Service of our Environment, Cannes, France, 4–6 November 1992 (ESA SP-359)*, pages 823–828, March 1993.

139. P. Wang, J. Chong, and H. Wang. Ship detection of the airborne SAR images. In *IEEE 2000 International Geoscience and Remote Sensing Symposium (IGARSS'00)*, volume 1, pages 348–50, 2000.

140. D. Werle. Radarsat SAR azimuth ambiguity patterns - the ghost fleet of Halifax Harbour and implications for applications. In *International Symposium, Geomatics in the Era of RADARSAT (GER'97)*, 1997. Ottawa, Canada, May 25-30.

141. M.L. Williams and Blucher G. Enhancing FOPEN performance at L-band using SAR polarimetry. In *5th European Conference on Synthetic Aperture Radar (EU-SAR2004)*, 2004.

142. M. Yeremy, J.W.M. Campbell, K. Mattar, and T. Potter. Ocean surveillance with polarimetric SAR. *Canadian Journal of Remote Sensing*, 27(4):328–344, 2001.

143. M. Yeremy, G. Geling, M. Rey, B. Plache, and M. Henschel. Results from the CRUSADE ship detection trial: polarimetric SAR. In *IEEE 2002 International Geoscience and Remote Sensing Symposium (IGARSS'02)*, 2002.

144. M.L. Yeremy, A. Beaudoin, J.D. Beaudoin, and G.M. Walter. Global shoreline mapping from an airborne polarimetric SAR: Assessment for RADARSAT 2 polarimetric modes. Technical Report DREO-TR-2001-056, Defence Research Establishment Ottawa, Canada, November 2001. 67 pages.

145. M.L. Yeremy, C. Livingstone, K. Mattar, L. Gallop, and J. Lang. Polarimetric inter-ferometry experimental trials for years 2001 and 2002: Experiment design, ground truthing, data quality and analysis. Technical Report DRDC Ottawa TM 2003-142, DRDC Ottawa, Canada, September 2003. 142 pages.

146. Yu Yongjian, Huang Shun-Ji, and A. Torre. Development of an automatic target detection and characterisation system in polarimetric SAR images. In *IEEE 1995 International Geoscience and Remote Sensing Symposium (IGARSS'95)*, volume 3, pages 1822–4, 1995.

The State-of-the-Art in Ship Detection in Synthetic Aperture Radar Imagery.

D. J. Crisp

Number of Copies

DEFENCE ORGANISATION

Task Sponsor

DGISREW	1

S&T Program

Chief Defence Scientist	
FAS Science Policy	
AS Science Corporate Management	1
Director General Science Policy Development	
Counsellor, Defence Science, London	Doc Data Sheet
Counsellor, Defence Science, Washington	Doc Data Sheet
Scientific Adviser to MRDC, Thailand	Doc Data Sheet
Scientific Adviser Joint	1
Navy Scientific Adviser	Doc Data Sheet and Dist List
Scientific Adviser, Army	Doc Data Sheet and Dist List
Air Force Scientific Adviser	1
Scientific Adviser to the DMO M&A	Doc Data Sheet and Dist List
Scientific Adviser to the DMO ELL	Doc Data Sheet and Dist List
Director of Trials	1

Information Sciences Laboratory

Chief, Intelligence, Surveillance and Reconnaissance Division	1
Research Leader, Imagery Systems	1
Head, Image Analysis & Exploitation	1
Guy Blucher	1
Dr David Booth	1
David I. Kettler	1
Dr Nicholas J. Redding	4
Merrilyn Fiebig	1
Dr Vittala Shettigara	1
Ray Oermann	1
Bob Whatmough	1

Head, Imaging Radar Systems	1
Dr David Crisp	1
Rodney Smith	1
Dr Mark Williams	1

DSTO Library and Archives

Library, Edinburgh	1 and Doc Data Sheet
Australian Archives	1

Capability Systems Division

Director General Maritime Development	Doc Data Sheet
Director General Land Development	1
Director General Aerospace Development	1
Director General Information Capability Development	Doc Data Sheet

Office of the Chief Information Officer

Deputy Chief Information Officer	Doc Data Sheet
Director General Information Policy and Plans	Doc Data Sheet
AS Information Structures and Futures	Doc Data Sheet
AS Information Architecture and Management	Doc Data Sheet
Director General Australian Defence Simulation Office	Doc Data Sheet

Strategy Group

Director General Military Strategy	Doc Data Sheet
Director General Preparedness	Doc Data Sheet

HQAST

SO (Science) ASJIC	Doc Data Sheet

Navy

SO (SCIENCE), COMAUSNAVSURFGRP, NSW	Doc Data Sheet and Dist List
Director General Navy Capability, Performance and Plans, Navy Headquarters	Doc Data Sheet
Director General Navy Strategic Policy and Futures, Navy Headquarters	Doc Data Sheet

Army

ABCA National Standardisation Officer, Land Warfare Development Sector, Puckapunyal	Doc Data Sheet (pdf format)
SO (Science), Deployable Joint Force Headquarters (DJFHQ)(L), Enoggera QLD	Doc Data Sheet
SO (Science), Land Headquarters (LHQ), Victoria Barracks, NSW	Doc Data Sheet and Exec Summ

Air Force

SO (Science), Headquarters Air Combat Group, RAAF Base, Williamtown	Doc Data Sheet and Exec Summ

Intelligence Program

DGSTA, Defence Intelligence Organisation	1
Manager, Information Centre, Defence Intelligence Organisation	1 (pdf format)
Assistant Secretary Corporate, Defence Imagery and Geospatial Organisation	1

Defence Materiel Organisation

Head Airborne Surveillance and Control	Doc Data Sheet
Head Aerospace Systems Division	Doc Data Sheet
Head Electronic Systems Division	Doc Data Sheet
Head Maritime Systems Division	Doc Data Sheet
Head Land Systems Division	Doc Data Sheet
Head Industry Division	Doc Data Sheet
Chief Joint Logistics Command	Doc Data Sheet
Management Information Systems Division	Doc Data Sheet
Head Materiel Finance	Doc Data Sheet

Defence Libraries

Library Manager, DLS-Canberra	Doc Data Sheet
Library Manager, DLS-Sydney West	Doc Data Sheet

UNIVERSITIES AND COLLEGES

Australian Defence Force Academy Library	1
Head of Aerospace and Mechanical Engineering, ADFA	1
Deakin University Library, Serials Section (M List), Geelong, Vic	1
Hargrave Library, Monash University	Doc Data Sheet
Librarian, Flinders University	1

OTHER ORGANISATIONS

National Library of Australia	1
NASA (Canberra)	1
Government Publications Librarian, State Library of New South Wales	Doc Data Sheet

INTERNATIONAL DEFENCE INFORMATION CENTRES

US Defense Technical Information Center	2

UK Defence Research Information Centre	2
Canada Defence Scientific Information Service	1 (pdf format)
NZ Defence Information Centre	1

ABSTRACTING AND INFORMATION ORGANISATIONS

Library, Chemical Abstracts Reference Service	1
Engineering Societies Library, US	1
Materials Information, Cambridge Scientific Abstracts, US	1
Documents Librarian, The Center for Research Libraries, US	1

INFORMATION EXCHANGE AGREEMENT PARTNERS

National Aerospace Laboratory, Japan	Doc Data Sheet
National Aerospace Laboratory, Netherlands	Doc Data Sheet

SPARES

DSTO Edinburgh Library	5

Total number of copies:	52

DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION DOCUMENT CONTROL DATA		1. CAVEAT/PRIVACY MARKING	
2. TITLE The State-of-the-Art in Ship Detection in Synthetic Aperture Radar Imagery.		3. SECURITY CLASSIFICATION Document (U) Title (U) Abstract (U)	
4. AUTHOR D. J. Crisp		5. CORPORATE AUTHOR Information Sciences Laboratory PO Box 1500 Edinburgh, South Australia, Australia 5111	

6a. DSTO NUMBER DSTO–RR–0272	6b. AR NUMBER 013–053		6c. TYPE OF REPORT Research Report	7. DOCUMENT DATE May, 2004
8. FILE NUMBER E9505 / 28 / 44	9. TASK NUMBER JTW 01/218	10. SPONSOR DGISREW	11. No OF PAGES 115	12. No OF REFS 146

13. URL OF ELECTRONIC VERSION http://www.dsto.defence.gov.au/corporate/ reports/DSTO–RR–0272.pdf	14. RELEASE AUTHORITY Chief, Intelligence, Surveillance and Reconnaissance Division

15. SECONDARY RELEASE STATEMENT OF THIS DOCUMENT

Approved For Public Release

OVERSEAS ENQUIRIES OUTSIDE STATED LIMITATIONS SHOULD BE REFERRED THROUGH DOCUMENT EXCHANGE, PO BOX 1500, EDINBURGH, SOUTH AUSTRALIA 5111

16. DELIBERATE ANNOUNCEMENT

No Limitations

17. CITATION IN OTHER DOCUMENTS

No Limitations

18. DEFTEST DESCRIPTORS

Algorithms
Image analysis
Surface ship detection
Synthetic aperture radar

19. ABSTRACT

This report is a review of the publicly available literature on algorithms for detecting ships in Synthetic Aperture Radar (SAR) imagery, with the aim of recommending algorithms for inclusion in the Analysts' Detection Support System (ADSS). The ADSS is a software system being developed at DSTO for the automatic detection of targets in SAR imagery. We outline the motivations, theory and justifications of the various approaches so that educated comparisons and evaluations can be made. Most current research on ship detection is based on detecting the ship itself rather than its wake. Detection of ship wakes is not addressed in any depth. Ship detection systems generally consist of several stages: land masking; preprocessing; prescreening; and discrimination. We consider each of these stages in turn and discuss the various algorithms which have been used. A basic system for low to medium resolution single channel SAR imagery is recommended and suggestions for improvements and alternatives are made.